Modeling insect inspired mechanisms of neural and behavioral plasticity

Inaugural dissertation
to obtain the academic degree
Doctor rerum naturalium (Dr. rer. nat.)
submitted to the department of biology, chemistry, and pharmacy
at the Freie Universität Berlin

by

Joachim Haenicke
from Bad Oeynhausen

May 2015

The research presented in this dissertation was carried out from October 2010 until May 2015 in the Theoretical Neuroscience & Neuroinformatics group at the Freie Universiät Berlin, under the supervision of Prof. Dr. Martin P. Nawrot.

1st Reviewer: Prof. Dr. Martin P. Nawrot
2nd Reviewer: Prof. Dr. Hans-Joachim Pflüger

Date of defense: 10.11.2015

This dissertation is based on the following three manuscripts:

Conditioned behavior in a robot controlled by a spiking neural network

Authors: Lovisa Irpa Helgadottir[1,2], Joachim Haenicke[1,2], Tim Landgraf[3], Raul Rojas[3], Martin P. Nawrot[1,2]

Author contributions: Network model by J.H. Technical implementation by L.I.H., J.H., and T.L. Manuscript by L.I.H., J.H., T.L., R.R., and M.P.N.

Acknowledgments: We would like to thank Michael Schmuker, Hamid Mobalegh and Jan Meyer for their help and contribution.

Manuscript status: The manuscript has been published.

6th International IEEE/EMBS Conference on Neural Engineering, NER (2013, pp. 891–894). http://dx.doi.org/doi:10.1109/NER.2013.6696078

Neural correlates of odor learning in the microglomerular circuitry of the honeybee mushroom body

Authors: Joachim Haenicke[1,2], Nobuhiro Yamagata[5], Hanna Zwaka[2], Martin P. Nawrot[1,4], Randolf Menzel[2]

Author contributions: J.H.: Data analysis, manuscript. N.Y.: Experiment, revision. H.Z.: Confocal scans, revision. M.P.N. and R.M.: Comments on data analysis, revision.

Manuscript status: A revised version of the current manuscript will be submitted for publication in an international peer reviewed journal.

Distributed plasticity in a network model of the honeybee brain explains behavioral responses in classical conditioning experiments

Authors: Joachim Haenicke[1,2], Evren Pamir[1], Martin P. Nawrot[1,4]

Author contributions: J.H.: Data collection and analysis, model design, software implementation, simulation, manuscript. E.P.: Concept, data collection and analysis, revision. M.P.N.: Data collection and analysis, revision.

Acknowledgments: We thank Michael Schmuker and Farzad Farkhooi for helpful comments on this work.

Manuscript status: A revised version of the current manuscript will be submitted for publication in an international peer reviewed journal.

Author affiliations:

1 Bernstein Center for Computational Neuroscience, Berlin.
2 Institute of Biology, Freie Universität Berlin.
3 Institute of Computer Science, Freie Universität Berlin.
4 Institute for Zoology, University of Cologne
5 Department of Developmental Biology and Neurosciences, Tohoku University

Zusammenfassung

Die vorliegende Arbeit untersucht die Verarbeitungsprozesse im Inektenhirn, die der Duftwahrnehmung und dem assoziativen Lernen zugrunde liegen. Kapitel 2 stellt eine vereinfachte Variante eines Modells der Duftverarbeitung im Fliegenhirn dar, welche über die Verarbeitung von sensorischen Daten hin zur Generierung motorischer Befehle die Aktivität eines Roboters steuert. Diese auf einem spikenden neuronalen Netzwerk beruhende Kontrollarchitektur wird im Rahmen eines einfachen Konditionierungsexperiments getestet. In Kapitel 3 wird die neuronale Aktivität an synaptischen Endigungen im Pilzkörpereingang des Bienenhirns untersucht. Hierfür wurden die im Rahmen eines Lernexperiments der klassischen Konditionierung von Bienen gewonnen Daten analysiert. Es wird gezeigt, dass individuelles Tierverhalten mit den lernbedingten Veränderungen in der neuronalen Aktivität hinsichtlich belohnter Düfte korreliert. Eine mögliche Erklärung hierfür wird in Kapitel 4 im Rahmen eines Netzwerkmodells des Bienenhirns detailliert beschrieben. Einzelne Verarbeitungsschritte der Duftwahrnehmung und -verarbeitung werden in einem rechnerbasierten Modell abstrahiert abgebildet und durch die von Kapitel 3 gestützte Hypothese zur neuronalen Plastizität ergänzt. Diese wird anschließend anhand einer Vielzahl von elementaren und nicht-elementaren Lernparadigmen evaluiert. Kapitel 4 stellt somit eine Verbindung zwischen den Verhaltensdaten und den neurophysiologischen Erkenntnissen zur Duftverarbeitung im Bienenhirn her.

Summary:

This thesis investigates neural activity underlying olfactory processing and associative learning in the insect brain. Chapter 2 presents a simplified version of a model of olfaction in the fly brain, that processes sensory input in order to generate appropriate motor commands that control the activity of a robot. This spiking neural network control architecture is tested in a simple conditioning experiment. Chapter 3 investigates the neural activity in synaptic terminals at the mushroom body input of the honeybee brain. For this, data from a classical conditioning experiment was analyzed. It is shown that individual behavior is correlated with learning-induces changes in neural responses towards rewarded odors. A possible explanation for this is described in detail in chapter 4 in form of a network model of the honeybee brain. Individual stages of olfactory processing are expressed in an abstract computation model, including a hypothesis on neural plasticity that is supported by the results of chapter 3. Subsequently, this hypothesis is evaluated based on a rich collection of data from elemental and non-elemental learning paradigms. Therefore, chapter 4 provides a link between behavior and neurophysiological knowledge about odor processing in the honeybee brain.

Keywords:
associative learning, autonomous robots, classical conditioning, computational modeling, honeybee, insect brain, neuronal computation, plasticity, spiking neural networks

Contents

Chapter 1

General introduction

1.1 The insect brain: An invaluable model system for neuroscience

The brains of insects are considerably smaller than those of vertebrates, yet they give rise to a range of complex behavioral capabilities such as highly differentiated motor repertoires, extensive social structures and cognition (Chittka and Niven, 2009). Among insects, fruit flies (*Drosophila melanogaster*) have become one, if not the most widely studied animal model, in science in general and in neuroscience in particular, a process that has been strongly shaped through the development of genetic tools (Weiner, 1999, for a brilliant bibliographic report on this subject). Although the genetic toolbox is considerably smaller for other insect species, several of these have become established animal models. Among them, eusocial insect families of the order hymenoptera, in particular ants and bees, take a special place, as they show remarkable skills of communication, navigation, and learning, and by this have become a source of inspiration for the design of autonomous robots and control algorithms (Arena et al., 2013).

While some of their strategies to solve such complicated tasks might turn out to be specific to their kind, others will provide insights into fundamental principles of biological information processing. One likely candidate for the latter is olfactory processing. Here, several structural and functional components are shared across phyla (Hildebrand and Shepherd, 1997; Ache and Young, 2005) and most likely have developed independently (Eisthen, 2002). The fact that the precision, speed, and general performance of chemosensation even in fruit flies is beyond reach for current technical implementations, such as artificial noses, can at least in part be explained by a lack of knowledge about the underlying principles.

This thesis is focused on internal processes of neural plasticity in honeybees (*Apis mellifera*) during olfactory conditioning experiments, and by this tries to contribute to a better understanding of olfactory processing and learning in insects. It takes three different approaches: Chapter 2 presents a robotic platform for the implementation and evaluation of spiking neural network

control architectures. Inspired by models of olfactory conditioning in the fruit fly (Wessnitzer et al., 2012) and honeybee (Haenicke et al., 2012), a simplified model of the insect sensory-to-motor network is combined with a biologically inspired synaptic learning rule and tested in an absolute conditioning experiment. A differential conditioning experiment is performed within the study presented in chapter 3, where neural activity in the honeybee brain was recorded before and after training in order to identify learning dependent plasticity. The target region of the corresponding study were microglomerular complexes of the mushroom body calyx, a candidate region to play an important role for learning and memory in the honeybee. Chapter 4 presents a network model of olfactory processing in the honeybee that was inspired in part by the outcome of this experiment and entails a plasticity mechanism in the mushroom body calyx. Within a computational framework, model predictions were calculated for conditioned behavior in a variety of elemental and non-elemental learning paradigms and are evaluated with respect to the corresponding honeybee data. In that study, I have tried to express the known physiological components of the honeybee brain that are relevant to odor processing during classical conditioning in form of an abstract computational model. Therefore, the work strongly relied on the multitude of studies that have investigated the physiology of honeybee olfaction with various methods.

1.2 Investigating the physiology of olfactory processing and learning in honeybees

Our image, i.e. understanding of the world is intimately linked to the methods we use to create images. The investigation of insect physiology has, thus, been defined by the development of imaging techniques, such as optical microscopy. Incidentally, the oldest published image known to have been made with a microscope was that of bees (Fig. 1.1). Crucial to visualizing detailed structures in organic tissue, though, was the development of effective staining methods, explaining the temporal gap between the invention of optical microscopy and the first detailed images of neural tissue. To name one important dye, haematoxylin is in use since more than a hundred years (Allison, 1999) and was also involved in early studies that characterized details of the neural architecture in the honeybee brain (Kenyon, 1896; Witthöft, 1967). Witthöft was able to stain the nuclei of neurons in the bee brain and reported their number for a female worker bee to amount to 851,458$\pm$15%.

A growing number of imaging techniques has allowed to identify the prominent neuropils involved in olfactory processing in insects. The antennal lobe (AL) is the first-order olfactory neuropil, that receives input from receptor neurons of the antenna (ANT). Projections from the AL reach two further neuropils, the mushroom body (MB) and the lateral horn (LH). Various methods have been applied to assess the functional properties of these structures, many of these with respect to learning and plasticity. A lack of appropriate techniques, though, has so far hindered a detailed investigation of the LH. Surprisingly, only a few studies exist that inves-

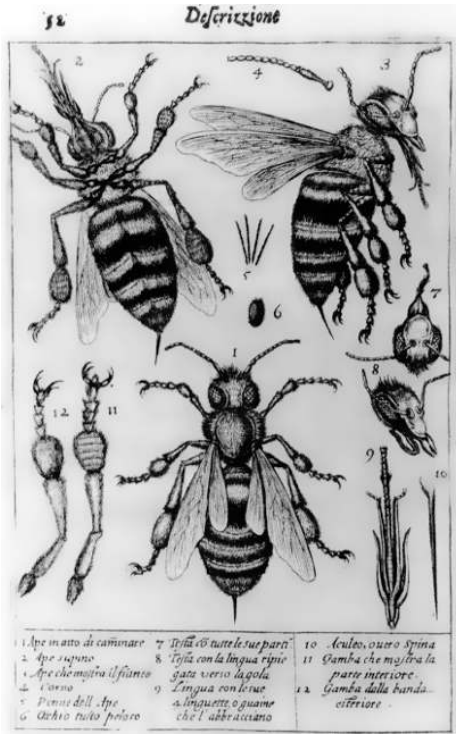

Figure 1.1: The oldest published image known to have been made with a microscope: bees by Francesco Stelluti, 1630. Image: public domain.

tigated details of receptor neurons in the honeybee ANT (Esslen and Kaissling, 1976; Lacher, 1964; Akers and Getz, 1993). However, both of the other two neuropils, AL and MB, have been numerously exposed to various techniques. To the exotic ones belong the analysis of current source density in the MB (Kaulen et al., 1984) and the cooling of localized areas of both AL and MB with small metal probes (Erber et al., 1980). Cellular details of synaptic interactions in the input-region of the MB have been investigated by electron microscopy (Ganeshina and Menzel, 2001), while, in a rather coarse approach, honeybee larvae were treated with hydroxyurea, causing varying degrees of MB ablations (Malun et al., 2002). Across AL and MB, a multitude of studies have used both intra- and extracellular recording techniques to describe the electrophysiological properties of the corresponding neuron types, though the most abundantly applied method is the recording of calcium-dependent neural activities in these structures.

Based on data from calcium-imaging, I show in chapter 3 that neural correlates of learning appeared in synaptic terminals at the input site of the MB. The possible sources of the observed plasticity are discussed in the light of learning-induced changes in neural activity that have been reported for different neuron types in both AL and MB. In addition, Chapter 4 offers a possible explanation in form of an abstract network model. With this, I tried to link internal computations in ANT, AL, MB, and LH to the behavioral plasticity observed in honeybees during classical conditioning. Prior to this, a simplified model version of the insect olfactory system that is integrated in a robotic platform is presented in the following chapter.

"It is certain that there may be extraordinary mental activity with an extremely small absolute mass of nervous matter: thus the wonderfully diversified instincts, mental powers, and affections of ants are notorious, yet their cerebral ganglia are not so large as the quarter of a small pin's head. Under this point of view, the brain of an ant is one of the most marvellous atoms of matter in the world, perhaps more so than the brain of a man." (Darwin, 1871)

Chapter 2

Conditioned behavior in a robot controlled by a spiking neural network

Abstract

Insects show a rich repertoire of goal-directed and adaptive behaviors that are still beyond the capabilities of today's artificial systems. Fast progress in our comprehension of the underlying neural computations make the insect a favorable model system for neurally inspired computing paradigms in autonomous robots. Here, we present a robotic platform designed for implementing and testing spiking neural network control architectures. We demonstrate a neuromorphic real-time approach to sensory processing, reward-based associative plasticity and behavioral control. This is inspired by the biological mechanisms underlying rapid associative learning and the formation of distributed memories in the insect.

Chapter 3

Neural correlates of odor learning in the microglomerular circuitry of the honeybee mushroom body

Plasticity in the mushroom body calyx correlates with learning success during classical conditioning of the honeybee.

Summary

In the honeybee, a rich body of evidence suggests the mushroom body (MB) to play a key role in memory formation and recall (Menzel, 2012). The exact localization of learning-related plasticity in the MB circuitry, however, remained elusive. Here, we studied neural plasticity in the MB calyx of the honeybee in relation to the learning success in individual animals. We performed classical conditioning experiments where a specific odor is paired with a sugar reward (Bitterman et al., 1983). Learning performance was monitored by the conditioned response (CR) behavior while we used Ca-imaging to measure the physiological odor responses in individual boutons in the MB calyx. These boutons form synaptic hubs between excitatory projection neurons (PNs) from the antennal lobe encoding the olfactory stimulus, the postsynaptic excitatory Kenyon cells, and GABAergic MB feedback neurons (Rybak and Menzel, 1993; Rybak and Menzel, 2010; Groh and Rössler, 2011). The octopaminergic reward pathway converges with this microcircuit (Sinakevitch et al., 2013). In the course of learning, individual boutons could show an overall increase or decrease of their Ca response, both for the rewarded stimulus (CS+) and the unrewarded control stimulus (CS-). The amount of neural plasticity induced for the CS+ was strongly and positively correlated with the learning success across individual animals. No significant correlation was observed for the changes in the CS- response. The temporal profile of the induced changes matched the Ca response dynamics of the inhibitory GABAergic feedback neurons. We hypothesize that the observed plasticity in the microglomerulus relates to the learned value of a stimulus, signifies short-term memory and underlies the conditioned response behavior.

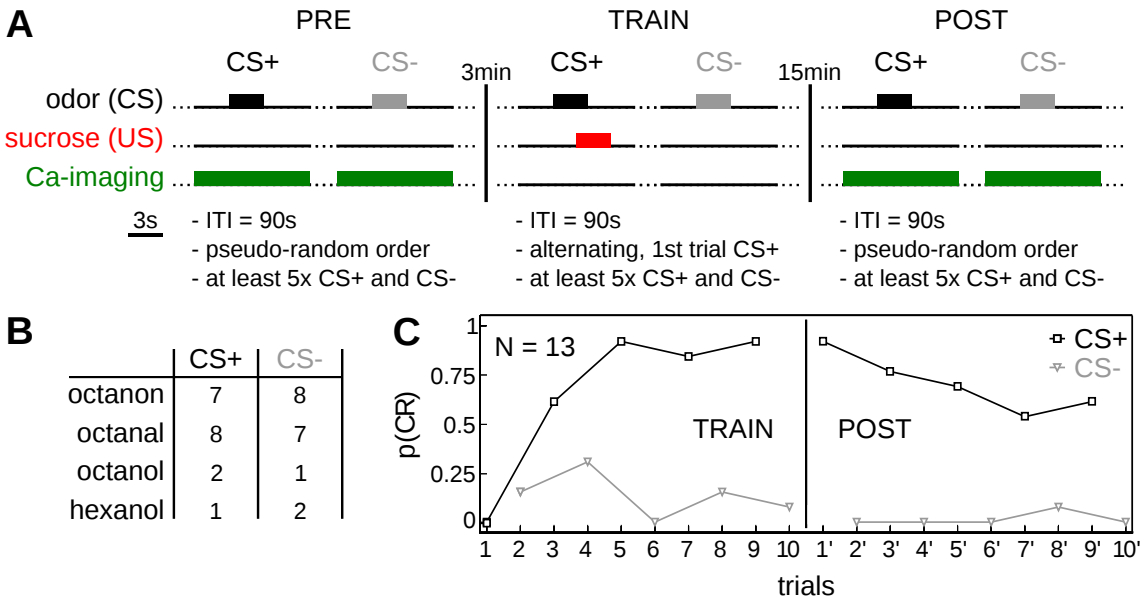

Figure 3.1: **Experimental paradigm and behavioral results.** Bees (N=18) were trained in a differential conditioning paradigm. (A) Differential conditioning. In the pre-training phase (PRE) two odors were applied in pseudo-random order, at least five times each. The subsequent training phase (TRAIN) always started with the sucrose-rewarded odor (CS+), alternating with the unrewarded odor (CS-). In the final post-training phase (POST) both odors were repeatedly presented again in a pseudo-random order. Ca-imaging was performed during phases PRE and POST. (B) Table of odor names and number of experiments in which they served as CS+ or CS-, respectively. (C) Behavioral data for the learner bees (N=13) during TRAIN (trials 1-10) and POST (trials 11-20). The probability p(CR) of observing a CR in the population rapidly increased and saturated after two pairings of CS+ and US. A small fraction of bees initially generalized towards CS-. Multiple unrewarded test trials during the POST phase lead to a gradual reduction of responses to the trained odor.

3.1 Results

Behavioral performance: Conditioned responses

We trained honeybees in a differential conditioning paradigm as depicted in Fig. 3.1A (c.f. Experimental procedures). We monitored individual animal behavior expressed in the proboscis extension response (PER) (Bitterman et al., 1983) during all three experimental phases. None of the bees showed a PER in any of the odor stimulation trials before training (PRE). Bees that did not show a behavioral response to any of the CS+ trials in the training phase (TRAIN) or the post training phase (POST) were classified as non-learners (Pamir et al., 2011; Pamir et al., 2014). For the learner bees a few trials were sufficient to develop a stable conditioned response (CR) to the rewarded odor (CS+) while they hardly responded to the unrewarded odor (CS-) as illustrated by the group-averaged response probabilities p(CR) in Fig. 3.1 C. An asymptotic level of $p(CR) \approx 0.9$ was reached already after two rewarded trials, which matches the generally observed rapid learning dynamics in honeybees during classical conditioning (Pamir et al., 2014).

Bouton activity: Odor-evoked response profiles

We studied learning induced physiological plasticity at the level of PN boutons in the MB calyx. To this end, PNs of the antennal lobe that project via the lateral antennal lobe tract (Fig. 2A) were dye-loaded with the calcium-sensitive dye Fura-dextran and the lysine fixable dye tetramethylrhodamine-dextran (Fig. 2B) that allowed to vizualise individual PN boutons in the MB calyx (Fig. 2C). Each single bouton forms the core of a microglomerular structure and is contacted by several postsynaptic KCs as well as by inhibitory feedback neurons (Sinakevitch et al., 2013). We imaged multiple trials of odor-evoked Ca-dependent responses before (PRE) and after (POST) training (Fig. 2E). The locations of individual Ca-active boutons were identified (Fig. 2D) based on the trial-averaged responses, and for all trials individually the time-resolved Ca activity of each bouton was extracted (c.f. Experimental procedures).

First, we compared Euclidean distances between Ca-response profiles in pairs of trials within the same experimental phase (PRE or POST) with distances between trials of the two different phases (PRE vs. POST). The results verified that activity patterns were significantly more similar within the same experimental phase than across phases (Supplementary Fig. 3.6 A; KS test, $p<0.0001$). For further analyses we computed the time-resolved trial-averaged bouton activities separately for each experimental phase and stimulus type (Fig. 3.2 F).

We observed diverse odor response patterns across animals, while the odor responses of individual boutons were often similar within individual animals showing rather homogeneous response profiles (Fig. 3.4). On account of this, the average pairwise Euclidean distances between individual trial-averaged bouton activities within animals were considerably and significantly smaller than those between animals (KS test, $p< 10^{-4}$, Fig. 3.6 B).

A more detailed characterization of odor response types is given in table 3.1 in the supplements. It lists fractions of all 1652 bouton activities during experimental phase PRE with respect to the presence of excitatory ON- and OFF-responses during the first second after odor onset and offset, respectively. For this period, we found virtually no inhibitory odor responses.

To analyze the temporal response dynamics we calculated time-resolved odor-evoked response profiles separately for excitatory and inhibitory odor responses. In the majority we observed excitatory responses with a phasic-tonic profile and low response latencies that closely matched previous Ca-imaging data recorded in PNs of the antennal lobe (Peele et al., 2006; Fig. 3.2 G) and the typical intracellular spike responses in PNs (Krofczik et al., 2008; Strube-Bloss et al., 2011; Brill et al., 2013). The average inhibitory response matched previous data recorded from the inhibitory MB feedback neurons (Haehnel and Menzel, 2010) (Fig. 3.2 H). Most excitatory odor responses started within the first two imaging frames after odor onset with an average latency of $0.22\,\mathrm{s}$ (Fig. 3.2 I), whereas the latencies of inhibitory responses were more widely distributed with an average of $\approx 1\,\mathrm{s}$.

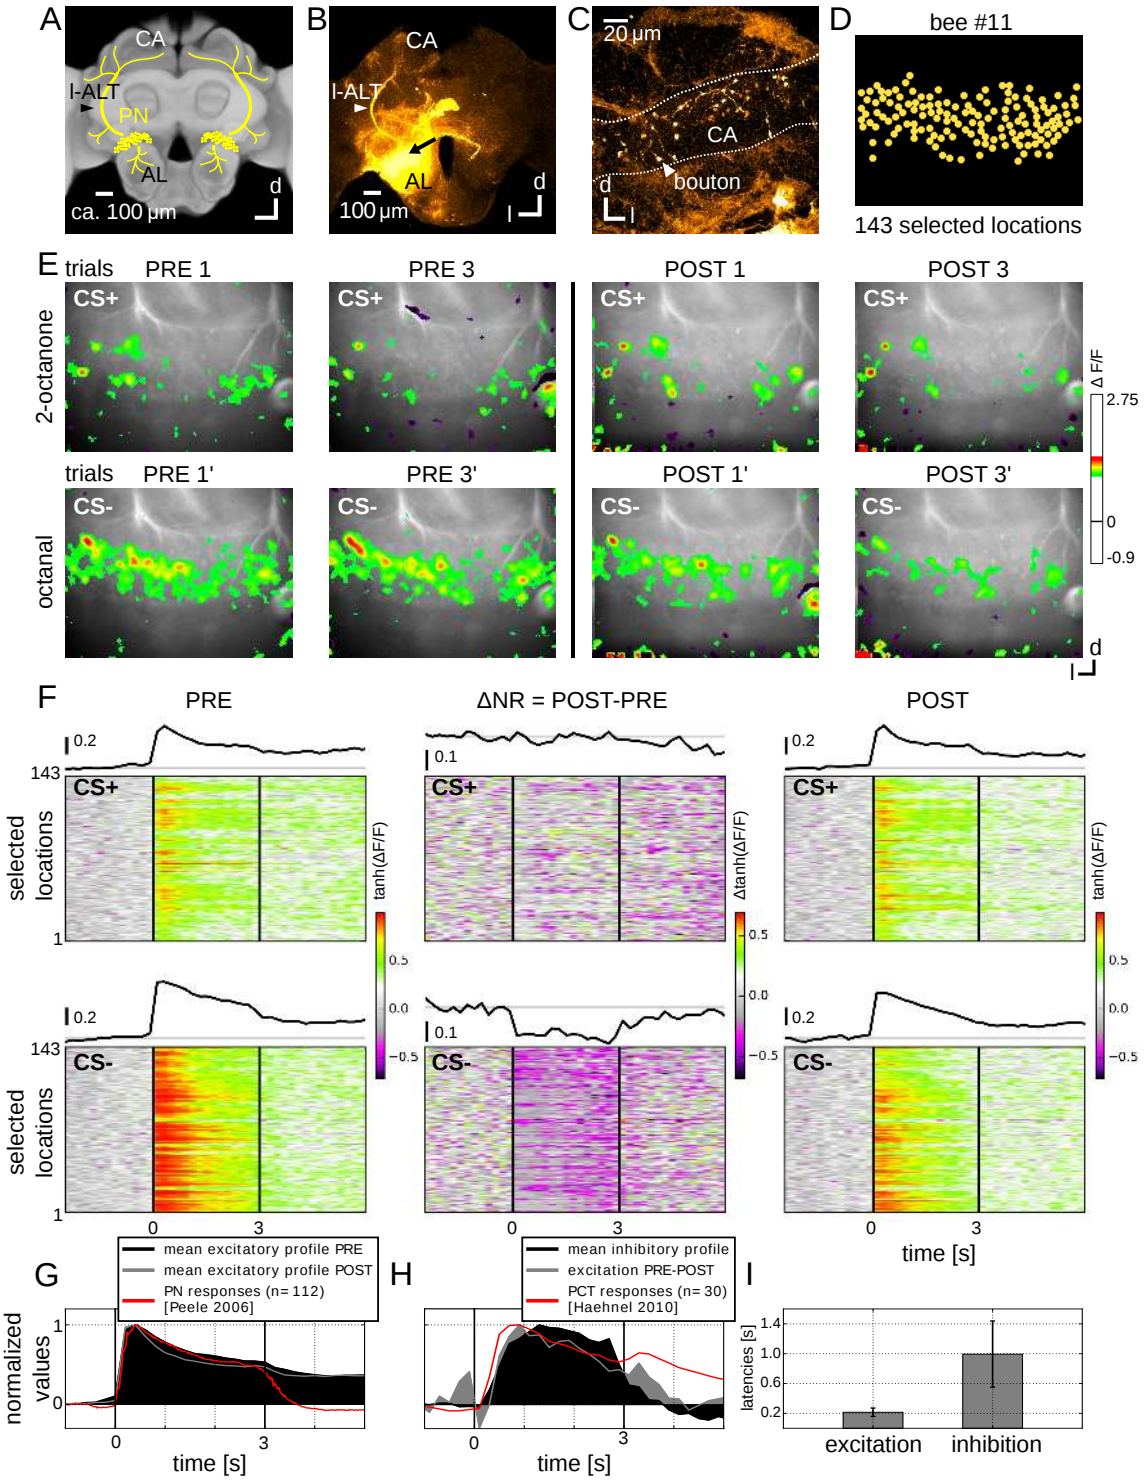

Figure 3.2: **Stainings and Ca-imaging from individual boutons.** (A) Frontal view of the honeybee brain with schematic overlay: Uniglomerular PNs (yellow) connect the antennal lobe (AL) with the mushroom body calyx (CA) via the lateral antennal lobe tract (l-ALT; indicated by triangle). (B) Confocal image of a stained brain, similar view as in A. PN somata were stained with a mixture of the Ca indicator Fura dextran and the neuro-tracer rhodamine dextran (sight of injection marked by black arrow). (C) Confocal image of rhodamine fluorescence shows synaptic boutons of PNs in the MB CA. (D) Identified locations of synaptic boutons for one of the animals. *Figure legend continues.*

Figure 3.2: *Figure legend continued.* (E) Color-coded Ca responses superimposed on raw fluorescence images for two trials (1 and 3) during the two experimental phases PRE (before training, left) and POST (after training, right) in response to the odors octanal (rewarded during training, i.e. CS+, upper row) and 2-octanone (CS-, lower row). Same animal as in D. (F) Temporal dynamics of Ca activity for all 143 identified boutons (color coded) and average trace for all boutons (black curve). Traces were calculated as the average of all trials within experimental phase PRE (left column) or POST (right column) for both CS+ (upper row) and CS- (lower row). Changes in the trial-averaged activity were calculated as the difference between POST and PRE (central column). (G) Comparison of the mean excitatory response profiles of experimental phases PRE (black) and POST (gray) with previously published PN Ca data (red, Peele et al., 2006). (H) Comparison of the mean inhibitory response profile (black) with the difference in excitatory response profiles (PRE-POST, gray) and previously published Ca data from MB inhibitory feedback neurons (red, Haehnel and Menzel, 2010). (I) Mean latencies and standard deviations of excitatory and inhibitory response profiles.

Bouton plasticity: Correlation with behavioral performance across individuals

We observed a variety of differences in bouton activity patterns between the PRE and POST phases (Fig. 3.5). Supplementary Table 3.1 lists fractions of all 1652 recorded boutons with respect to the presence of significant positive and negative changes in odor-evoked responses during the first second after odor onset and offset.

We quantified neural plasticity for each animal and stimulus type (CS+ / CS-) separately computing ΔNR as the absolute change in the Ca-dependent odor response in a single animal averaged across all boutons (see Experimental procedures). Likewise, we quantified the behavioral performance ΔCR of each individual bee during the training (TRAIN) and after the training (POST) by computing the fraction of CS+ trials that lead to a conditioned response (CR) and subtracting the fraction of responded CS- trials (see Experimental procedures). A high performance in the TRAIN phase implied a high performance in the POST phase (Pamir et al., 2014) with a significant correlation ($\rho = 0.61$, p<0.007).

Next, we tested whether the physiological plasticity observed at the level of individual PN boutons in the calyx correlates with the behavioral performance during learning (TRAIN). We find a clear and highly significant positive correlation between CS+ plasticity and behavioral performance (Fig. 3.3 A, $\rho = 0.76$, p<0.0002). This correlation is even more pronounced for the subset of animals that showed Ca responses with a high signal to noise ratio (SNR, see Experimental procedures, $\rho = 0.84$, p<0.0003). The training-induced neuronal plasticity in response to the CS- odor, however, did not show a correlation with the behavioral performance (Fig. 3.3 B), neither when considering all animals nor for the subgroup of animals with high SNR.

We confirmed the robustness of our correlation results using a jackknife approach where we repeatedly used all but one animal to test for positive correlations between neuronal plasticity and behavioral performance (Fig. 3.3 C,D). For CS+ the correlation is narrowly distributed around the above stated values in both groups, and the corresponding p-values suggest statistical significance in all cases. In contrast, the correlations for the CS- are low and not significant

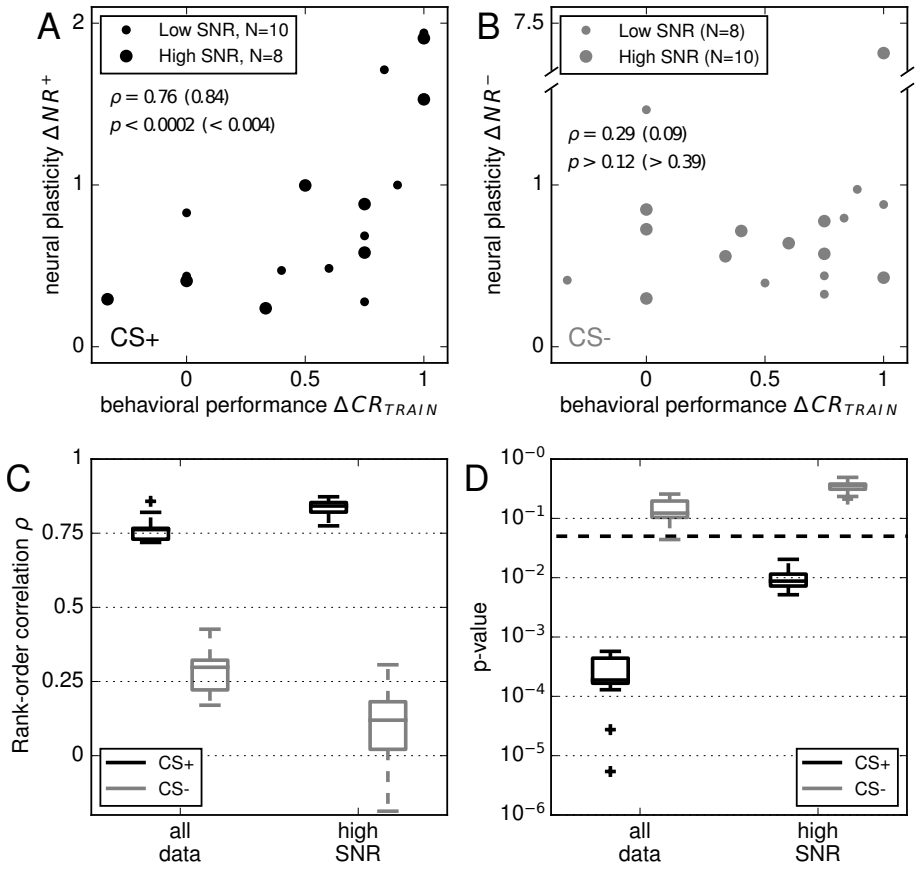

Figure 3.3: **Positive correlation between CR behavior and bouton plasticity.** (A) Scatter plot of behavioral performance during experimental phase TRAIN ΔCR^{TRAIN} against neural plasticity ΔNR^{+} for the rewarded odor. A subset of animals (large symbols) was selected on the basis of their high SNR. The correlation is highly significant. (B) ΔCR^{TRAIN} versus neural plasticity of unrewarded odors ΔNR^{-}. No significant correlation. (C) Jackknife resampling of Spearman's rank order correlations to estimate robustness and dependence on outliers. (D) Corresponding p-values. Thick dashed line indicates a level of significance of 0.05.

with the exception of one single iteration.

To asses the plasticity effect on the temporal odor response dynamics we calculated the difference in the trial-average excitatory odor response profile between the phases PRE and POST (Fig. 3.2 H). Interestingly, the time course of this difference curve resembles the odor response profiles of mushroom body extrinsic feedback neurons previously reported in (Haehnel and Menzel, 2010). This suggests that the observed plasticity in the odor responses at the level of boutons is due to learning-induced alterations of the inhibitory feedback response profile.

3.2 Discussion

Expression of short-term memory

Due to technical limitations, Ca responses could not be recorded during acquisition. However, neural response was recorded only 15 min after learning. For the rewarded odor, the amount of neural plasticity ΔNR^{+} (Fig. 3.3) was significantly correlated with the observed behavioral learning performance during acquisition ΔCR^{TRAIN}. We therefore hypothesize that the observed MB plasticity signifies short-term memory and that it is required for the generation of the CR behavior. Our results her are thus in contrast to the plasticity reported in MB output neurons that appeared only 3h after learning but not during acquisition which was interpreted as a signum of long term memory (Strube-Bloss et al., 2011).

Origin of the observed physiological plasticity

Previously described correlates of associative odor learning in honeybees that appeared in the neural activity of PNs have been primarily studied and interpreted on the basis of the AL network and for time scales related to long-term memory (Fernandez et al., 2009; Hourcade et al., 2009; Hourcade et al., 2010; Rath et al., 2011; Locatelli et al., 2012). An alternative hypothesis (Szyszka et al., 2008; Haehnel and Menzel, 2010) suggested associative plasticity to reside in the MB network, conveyed through a modulation of recurrent GABAergic signals of the MB feedback neurons. A possible mechanism could be disinhibition of the rewarded odor-evoked activity patterns at the level of individual boutons, as well as an increased inhibition of unrewarded activity patterns. Multiple physiological findings support this hypothesis: In isolated bee heads, odor responses of MB feedback neurons have been found to be modulated on account of odor conditioning after a single odor-reward pairing (Grünewald, 1999). Extracellular recordings from these MB feedback neurons showed their plasticity in classical conditioning experiments with olfactory and visual stimuli (Klinke, 2011) . The recent study by (Sinakevitch et al., 2013) described octopamine receptors to be co-localized with projections of the GABAergic feedback neurons that terminate in the MB lip-region. Ongoing research investigates associative learning with inhibitory feedback plasticity in a computational model of the honeybee MB.

In our data, for both CS+ and CS-, we observed a dominant decrease in the grand average of odor-evoked activity profiles during odor stimulation from experimental phase PRE to POST. The difference (PRE-POST) in these average excitatory profiles is closely related to the mean temporal structure of activities that we classified as inhibitory, and both resemble those of odor-evoked responses of PCT-neurons during odor stimulation (Haehnel and Menzel, 2010) (figure 3.2H). Taken together, these observations suggest that our recordings represent compound signals that were predominantly shaped by the neural activity of PNs and modulated by GABAergic feedback of PCT neurons.

Response similarities within animals

Within each of the animals, the temporal structures of individual bouton responses were often closely related to each other (compare figure 3.6 B). This result might be explained by the following reasoning. Due to technical limitations, we observed a specific, locally restricted patch of the MB calyx of each animal, comprising an estimated 2% of the total area. We may thus assume that this small section contained boutons from few and possibly similarly tuned PNs. PNs form multiple boutons along their axons (Yamagata et al., 2009) so that we likely observe correlated activity in a number of boutons. Recent evidence in the fruit fly suggests that projection neurons with similar tuning properties have a tendency to converge onto the same Kenyon cells (Gruntman, Turner 2013).

Effect of odor identities

In a majority of recordings, we applied the odorants octanal and octanon as conditioned stimuli. A systematic feature of these recordings is the variation in signal strength, with octanal eliciting mostly strong odor responses and octanon eliciting mostly weak odor responses, as illustrated by the bimodal distribution of signal qualities in figure 3.6C. Both odors were equally used as rewarded and as unrewarded stimulus, and we did not find any significant difference in behavioral performances with respect to odor identity. Therefore, we do not assume any direct effect on the observed relations between calcium-imaging data and behavioral data.

Physiological plasticity might be underestimated

While we observed animals with a low plasticity score ΔNR^+ that also showed a stable behavioral performance ΔCR, we did not observe animals with a high plasticity score accompanied by a weak behavioral performance (figure 3.3A). A likely explanation for this observation is that we capture only a fraction of neural changes that appeared in a specific animal due to the undersampling of a small fraction of the calyx. In addition, the staining procedure may have resulted in a varying amount of backfilled fibers/cells. Presumably through these and other related issues, the quality of odor-evoked signals differed across preparations and for different odors, and physiological changes may have been concealed. These reasons can lead to a reduced ΔNR^+ but not to an overestimated ΔNR^+, which could explain the observation in Fig. 3.3 A.

3.3 Experimental procedures

Preparation and dye loading

Bees were prepared as described in Yamagata et al., (2009). In short, foraging worker bees were collected, chilled and fixed in recording chambers with wax. The head capsule was opened and a mixture of the calcium-sensitive dye Fura-dextran (10 000 MW, Molecular Probes, Eugene, OR, USA) and the lysine fixable dye tetramethylrhodamine-dextran (10 000 MW, Molecular Probes, Eugene, OR, USA) was injected into the brain aiming for the soma cluster of the projection neurons (PNs) of the lateral atennal lobe tract (l-ALT). Then the head capsule was closed and the bees were fed until satiation and kept at 17–20°C for 8–24 h. Before measurements, the legs and wings were cut and the abdomen, thorax and mandibles were immobilized with wax. The antennae were fixed with n-eicosane and the calyces of the MB were exposed for measurements. Kwik-Sil Adhesive (World Precision Instruments, Inc.) was poured into the head capsule to completely stabilize the brain. After sealing the gaps with vaseline, the recording chambers were filled with Ringer solution (in mM: 130 NaCl, 7 CaCl2, 6 KCl, 2 MgCl2, 160 sucrose, 25 glucose and 10 HEPES, pH 6.7, 500 mosmol).

Calcium imaging

Calcium measurements were performed at room temperature with a sampling rate of 5 Hz, using a TILL-Photonics imaging setup mounted on a fluorescence microscope (Zeiss Axioskop, Germany). Fura was alternately excited at 340 and 380 nm. Exposure times were 15 and 60 ms, respectively. Each measurement started 3 s prior to stimulus onset and lasted for 10 s. Images were acquired through a 60×/0.9 NA water dipping objective (Olympus, Tokyo, Japan), a 410-nm dichroic mirror and a 440 nm long pass filter with an Imago CCD camera (640 × 480 pixels, 4 × binned on chip to 160 × 120). Pixel size was 1.47 × 1.47 μm, which allowed a resolution of single boutons of PNs.

Odor stimulation and conditioning

In the majority of bees, 2-octanone and octanal (Sigma, Deisenhofen, Germany) diluted to 10^{-2} were used. In three animals, 1-hexanol and 2-octanol (Sigma, Deisenhofen, Germany) were used, in two of them without dilution and in one of them diluted to 10^{-2} and 5x10^{-2}, respectively. The bees were exposed to a constant air stream. Injection of the odorant (40 μl soaked with 1 cm2 x 2 filter paper) into the constant air stream was switched on and off by a computer-controlled solenoid valve (Galizia et al., 1997). Odors were presented for 3 s at an inter-stimulus interval of 90 s.

Initially, sugar responses of the bees were checked and only responding animals were used further. Bees were then moved to the recording site and given some minutes to rest. The protocol for the classical conditioning experiment (Bitterman et al., 1983) was designed as follows (Fig.

1A). In the pre-training phase (PRE), each bee was exposed to two different odors. Both odors were presented repeatedly during 6-8 trials in a pseudo-random order. After an interval of 3 min the training phase (TRAIN) started. Bees were conditioned to one of the odors (CS+) by pairing the odor with an unconditioned reward stimulus (US) consisting of a drop of 30% sucrose that was presented with a delay of 2 s and lasted for 3 s. Conditioning always started with the rewarded odor (CS+) and the non-rewarded control odor (CS-) was alternately presented (5-10 trials). After a 15 min interval the bees were again exposed to both odors during at least 5 trials (5-8 trials) in a pseudo-random order (POST).

In all three experimental phases we monitored the animals' conditioned response (CR) behavior as expressed in the proboscis extension response (PER) by visual inspection. At the end of each experiment the sugar response of each bee was tested and only responding animals were included in the analyses. We performed calcium imaging from the MB calyx during the PRE and the POST phase simultaneously with the CS and US stimulation. Honeybees were selected for imaging during the POST phase according to their behavioral performance during training. Consequently, only a small fraction of non-learners were imaged during the POST phase.

Confocal microscopy

After Ca2+ measurements, the brain was dissected and fixed in 4% formaldehyde in Millonig's buffer overnight at 4°C. The brain was then rinsed in saline, dehydrated in ethanol, cleared in methyl salicylate, set into a chamber filled with methyl salicylate and observed with a confocal microscope (Leica TCS SP2; Leica, Wetzlar, Germany). The excitation wavelength was 543 nm with a Green HeNe laser. The entire brain was scanned with a 10×/0.4 NA air objective (Olympus, Tokyo, Japan). Where necessary, the AL was scanned with a 20×/0.70 NA air objective (Olympus, Tokyo, Japan) and the MB calyx was scanned with a 63×/1.32–0.6 NA oil objective (Olympus, Tokyo, Japan). Morphological images were acquired as an averaged raw fluorescence image of 380 nm during the measurements and was later unsharp mask-filtered in Photoshop (Adobe).

Behavioral responses

For experimental phases TRAIN and POST we characterized the corresponding behavioral performances ΔCR^{TRAIN} and ΔCR^{POST} for each individual animal by computing the difference between its PER-activity in all but the first CS+ trials CR_t^{CS+} and in the same number of CS- trials CR_t^{CS-} devided by the number of trials N_t. Thus, ΔCR is the difference between the

empirical probabilities of a conditioned response $p(CR)$ in CS+ and CS- trials.

$$\begin{aligned}\Delta CR &= \frac{1}{N_t}\sum_{t=2}^{N_{t+1}} CR_t^{CS+} - \frac{1}{N_t}\sum_{t=1}^{N_t} CR_t^{CS-} \\ &= p(CR^{CS+}) - p(CR^{CS-}), \Delta CR \in \Re^{(-1,1)}\end{aligned} \tag{3.1}$$

For the training phase (TRAIN), we excluded the first CS+ trial from this calculation because only the subsequent trials of the training can be considered as both acquisition and test trials. We defined bees as non-learners if they did not show a CR in any of the CS+ trials during both phases TRAIN and POST (Pamir et al., 2011).

Bouton activity: Odor-evoked response profiles

Recorded videos of calcium responses were preprocessed in IDL (RSI, Boulder, CO, USA) using custom scripts as described in (Yamagata et al., 2009). A mean of 15 frames during a single odor stimulation was calculated and displayed as a false-color image (Figure 2 E). A spatial low-pass filter (5 $\times$ 5 pixels) was applied to these images for better visualization. Individual boutons were determined as isolated activity spots (21 pixels) in the false-color images (Yamagata et al., 2009). We identified between 45 and 143 boutons per bee. For each bouton a response trace was calculated by averaging the signal of an activity patch without any filtering and correction.

All of the recorded bouton activity traces were included in the following analysis. Data was exported to Python and all subsequent analyses steps were performed using standard libraries for numerical/scientific programming in Python (NumPy, SciPy).

Overall, the observed calcium-dependent odor responses were dominated by excitatory signals and, thus, followed a skewed distribution with excitatory signals being overrepresented. This is mainly due to the fact that the inhibitory effect on the Ca response is bounded. It can only suppress the relatively small amount of spontaneous Ca activity (i.e. before odor stimulus) while an excitatory stimulation can lead to an arbitrary large increase in the Ca signal. Therefore, before analyzing odor response profiles and their training-specific changes in more detail, we transformed all data using the hyperbolic tangent function with the effect of gaining a better resolution of the inhibitory effect. This step facilitated the application of symmetric criteria for the classification of odor response profiles and their changes.

To acquire an overview of the different types of odor ON- and OFF-responses in the untrained animal (PRE) we first defined response intervals of 1s length starting with odor on- and offset, respectively. We then detected significant excitatory and inhibitory responses for each bouton based on the trial-averaged Ca-activity. During a baseline period of ($2.6s$ before stimulus onset we computed the mean a_0 and standard deviation σ_{a_0} across time. Excitatory/ inhibitory

responses were detected if the Ca activity exceeded $a_0 \pm 2.5\sigma_{a_0}$ in at least three of the five response samples that were recorded during a response interval. A threshold of $2.5\sigma_{a_0}$ proved to be suitable for separating the noisy signals that appeared during spontaneous activity from those values that are likely related to odor stimulation. In total, there are 9 possible combinations of odor ON- and OFF-response per bouton both for the rewarded (CS+) and non-rewarded odors (CS-).

We further computed the average dynamic response profiles and the response latencies (Figure 3.2G, H and I). For the average excitatory response we included all boutons that satisfied the condition for an ON-response. In order to capture the fewer, weaker and mostly delayed inhibitory ON-responses we extended the number of considered samples, requiring sub-threshold values in at least three of the 15 response samples that were recorded during the 3s odor presentation. For each of these excitatory and/or inhibitory responses, we calculated the response latency as the time corresponding to the first sample to exceed the defined threshold.

Bouton plasticity: Changes in odor response profiles

To provide a coarse overview of the changes in odor response profiles that exist between phases PRE and POST we defined a number of classes of changes in odor response on the basis of the change in activity during and shortly after odor stimulation (see 3.1). For this, we calculated the average c_0 and standard deviation σ_{c_0} across time for each bouton during baseline activity (prior to odor stimulation) with respect to odor type (CS+, CS-). We defined positive and negative changes as those that exceeded $c_0 + / - 2.5\sigma_{c_0}$ in at least one of the 20 samples that were recorded during the plasticity interval that we defined as the time between odor onset and 1s after odor offset. This threshold provided a suitable compromise for filtering most of the changes that appeared during spontaneous activity while keeping those that might be related to odor stimulation. Overall, we defined a set of four types of odor response changes for both the rewarded and the unrewarded odors (see table 3.1).

To further analyze the changes in average odor response profiles between phases PRE and POST, we assigned a single measure of neural plasticity ΔNR to each animal and stimulus type (ΔNR^+ for CS+ and ΔNR^- for CS-). For this, we computed the sum of the absolute change across all samples that exceeded $c_0 \pm 2.5\sigma_{c_0}$ in the plasticity interval on a per bouton basis. In a second step we averaged across boutons to quantify neural plasticity for each individual animal and each stimulus type.

Signal quality

We observed pronounced odor-dependent variations in response strength and signal quality. We analyzed this relationship by characterizing the signal quality of each animal as its average signal-

to-noise ratio (SNR) calculated according to equations 3.2 and 3.2, with $a_b(t)$ representing the time-varying activity trace of bouton b for a total of N_b boutons per animal. $\mu()$ and $\sigma^2()$ indicate the mean and the variance, respectively:

$$SNR_b = \left(\frac{P_{signal}}{P_{noise}}\right) = \frac{\mu\left(a_b(t)|_{t=0s}^{4s}\right)}{\sigma^2\left(a_b(t)|_{t=-2.6s}^{0s}\right)} SNR = \mu(SNR_b|_{b=1}^{N_b}) \tag{3.2}$$

We calculated SNR for each animal and stimulus type as an average of responses during both TRAIN and POST. Based on the bimodal distribution of these values we defined an empirical threshold to distinguish between weak (low SNR) and strong (high SNR) signals (see figure 3.6B for details).

Correlation analysis

We calculated Spearman's rank order correlation and the corresponding p-value as a test for positive linear correlations between measures of behavioral performances (ΔCR^{TRAIN}, ΔCR^{POST}) and neural plasticity (ΔNR^{+}, ΔNR^{-}). We performed correlation analyses separately for the full data set consisting of all 18 animals and for the subset of animals with recordings that showed strong signals according to the definition given in the previous paragraph (8 animals in the case of CS+ and 10 animals for CS-; see figure 3.6B for details). Because of the low number of samples, we applied jackknife resampling as an estimate for the robustness and dependence on outliers of the correlation results.

3.4 Supplemental information

	types of odor responses				types of response changes			
Stimulus	ON	OFF	both	none	up	down	both	none
CS+	39.5	3.3	9.9	47.2	14.2	33.0	7.8	45.0
CS-	40.9	2.7	12.3	43.6	18.7	31.4	5.7	44.2
both	15.8	0.0	0.8	16.6	2.8	11.0	0.6	20.7

Table 3.1: **Percentages of bouton types classified according to odor responses and response changes.**

Odor responses

Allthough boutons were visually identified as locations of odor-dependent activity, the activities of approx. 17% of boutons appeared too sparse or too noisy to satisfy our rather conservative criteria for an odor-evoked response (none+both). The fractions of bouton response profiles that were accepted as excitatory ON- and/or OFF-responses are roughly equally distributed in terms of CS+ and CS-, with around 40% that showed excitatory ON-responses, approx. 10% that showed excitatory ON- and OFF-responses, and fewer than 5% that showed only excitatory OFF-responses. Nearly all excitatory OFF-responses were exclusive to one of the applied stimuli, meaning that virtually none of the boutons showed excitatory OFF-responses for both CS+ and CS-.

More than half of the boutons (approx. 58%) responded to one of the odors only. This high odor specificity is mainly caused by the fact that we recorded weak or no responses at all to one of the two applied odors for nearly all of the animals (see figure 3.6B for details). This clear dichotomy of odor response profiles is mainly connected to stimulations with 2-octanone and octanal (about equally distributed between CS+ and CS-, compare figure 3.1B) in 15 out of 18 animals. 2-Octanone elicited only weak responses in 13 of these. On the contrary, octanal elicited rather pronounced signals in all but one of the recordings.

Response changes

Nearly 30% of all boutons did not show a significant change in their average response profile from PRE to POST. For the remaining 70% of boutons we report a number of salient features: The majority (approx. 80%) of boutons that changed their response profiles significantly did this in one direction only within the analyzed response windows, i.e. they either reduced or increased their odor response from PRE to POST. A minority of less than 13% of boutons showed signified changes in both directions in their dynamic response profile. Response reduction, i.e. an inhibition of the previous odor response, appeared more often (in more than 30% of the

cases) than response increase. Across all classes, the number of boutons are roughly equally distributed with respect to stimulus type (CS+/CS-). Hence, looking at the sign of the changes in odor response profiles did not reveal any apparent contrast between the rewarded and the unrewarded odor. A majority of the boutons showed specificity to either one of the odors in the changes. Just as for odor-evoked responses, this is related to the bimodal distribution of odor-dependent signal qualities (figure 3.6C).

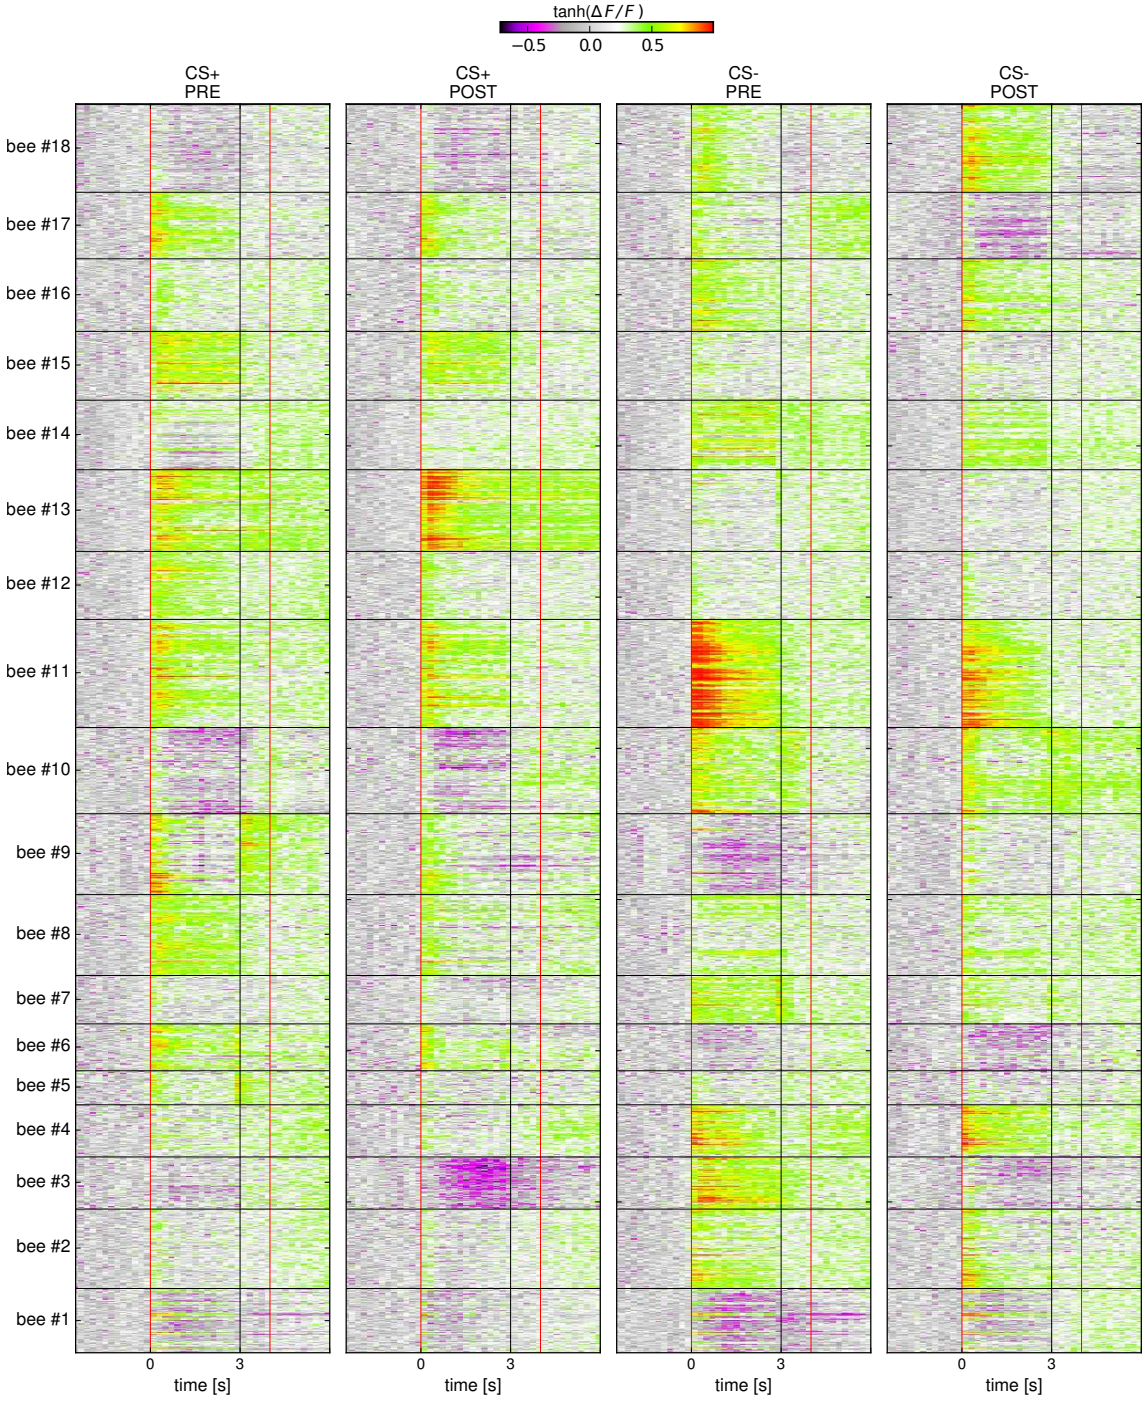

Figure 3.4: **Trial-averaged odor-evoked bouton activity patterns.** Color-coded calcium-dependent odor response profiles of all 18 bees, comprising traces of 1652 PN boutons recorded at the MB input for stimulations with CS+ (left two columns) and CS- (right two columns), respectively. For each bouton, response profiles were calculated as an average of all trials within each of the two experimental phases PRE (first and third column) and POST (second and fourth column), respectively. Black horizontal lines separate boutons of different animals, whose IDs are given as ordinate labels. Odors were applied between 0s and 3s. Measures of neural plasticity and SNRs were calculated on the basis of all time steps between 0s and 4s (enclosed by red horizontal lines).

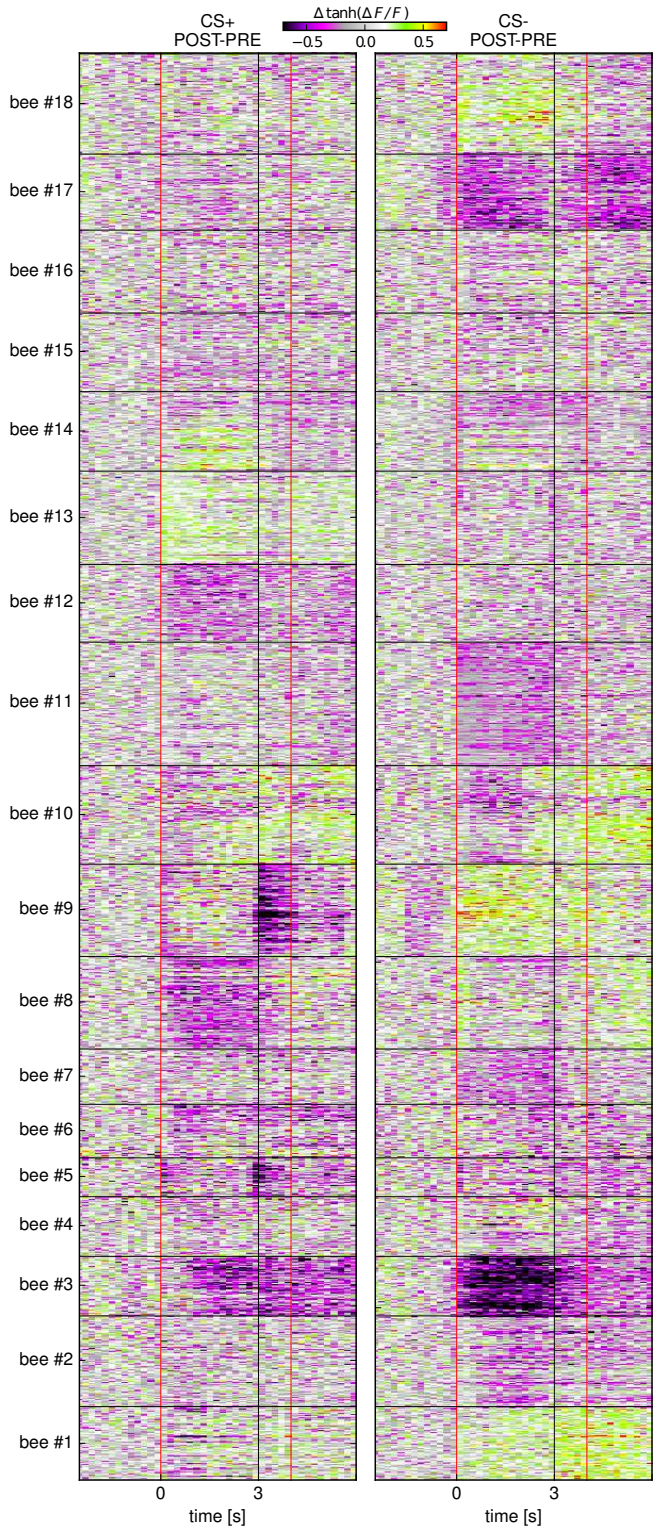

Figure 3.5: **Changes in trial-averaged odor-evoked bouton activity patterns.** Color-coded changes in calcium-dependent odor response profiles of all 18 bees, comprising traces of activity modulations of 1652 PN boutons recorded at the MB input for stimulations with CS+ (left column) and CS- (right column), respectively. Differences were calculated by subtracting averages of all trials of experimental phase PRE from those of experimental phase POST for each bouton. Black horizontal lines1 separate boutons of different animals, whose IDs are given as ordinate labels. Odors were applied between 0s and 3s. Measures of neural plasticity and SNRs were calculated on the basis of all time steps between 0s and 4s (enclosed by red horizontal lines).

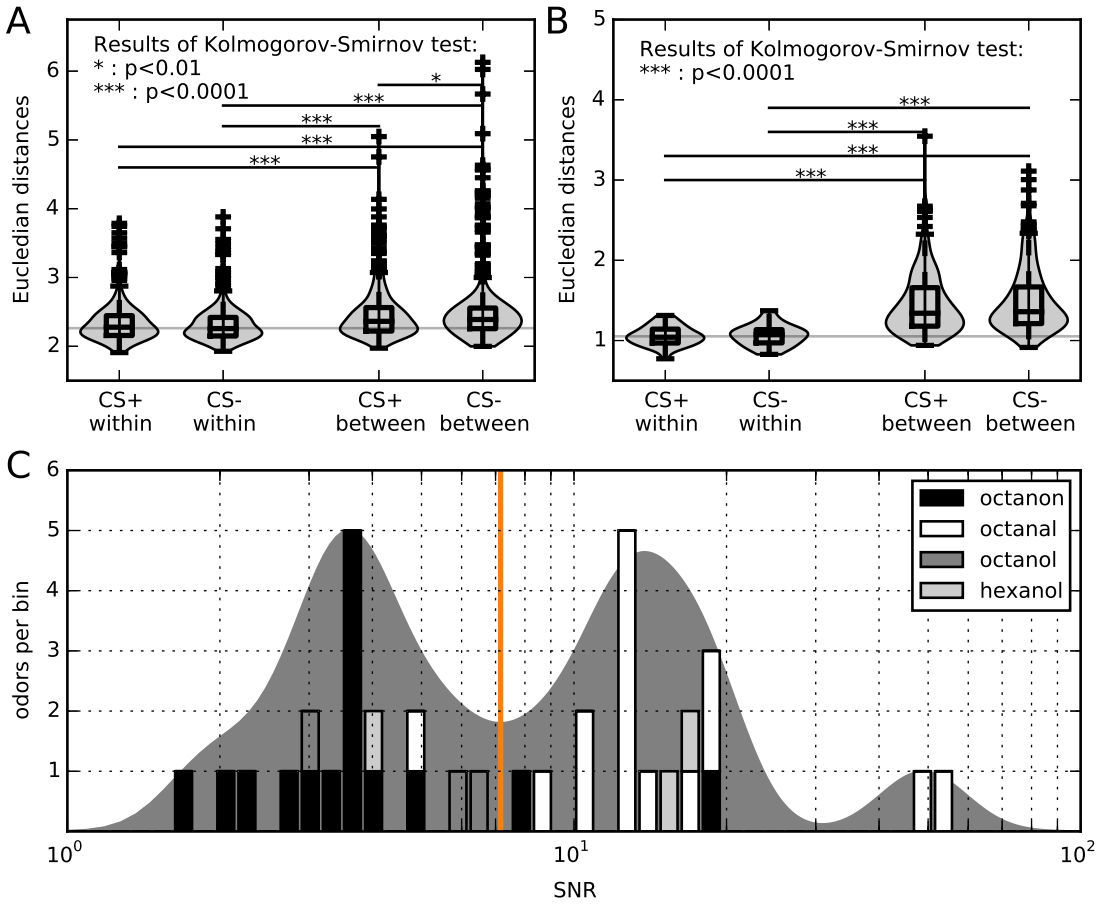

Figure 3.6: **Additional data properties.** A: Distributions of pairwise Euclidean distances of trials within and between different experimental phases. On average, bouton activity patterns are more similar to those of the same experimental phase than to those recorded during the other phase: Gray areas show smoothed distribution densities of Euclidean distances that were calculated as pairwise distances of all combinations of trials within or between experimental phases PRE and POST, averaged across boutons of each animal. Black boxplots show medians within boxes extending from lower to upper quartile, whiskers extend to the most extreme data points within 75%-25% of the inner quartile range, fliers represent data points that extend beyond the whiskers. Pairwise Kolmogorov-Smirnov statistics were applied as a two-sided test for the null hypothesis that two independent samples are drawn from the same continuous distribution. B: Distributions of average pairwise Euclidean distances between trial-averaged bouton activities within the same animal and between different animals. On average, bouton activity patterns are more similar to those of the same animal than to those recorded in other animals. C: Signal-to-noise ratios sorted according to odor type. Histogram of signal-to-noise ratios (SNR) with odor types represented in different shades of gray. For illustrative purposes, a Gaussian kernel density estimation of the underlying distribution was added as gray area in the background. The orange line indicates the median of all SNRs that was used as a decision threshold to separate strong from weak recordings. Applied odor combinations were either octanon and octanal or octanol and hexanol, with both pairs of odors distributed randomly between CS+ and CS- (compare with table in figure 3.1B).

Chapter 4

Distributed plasticity in a network model of the honeybee brain explains behavioral responses of classical conditioning experiments.

Abstract

The honeybee is a prominent model for studying the neural mechanisms underlying the formation and retrieval of associative memories. Here, we present a biologically motivated neural network model of the honeybee (Apis mellifera) brain that forms associative memories. In our network simulations we can reproduce the observed conditioned response (CR) behavior of honeybees in a variety of classical olfactory conditioning protocols. Our network model comprises peripheral olfactory receptor neurons, the antennal lobe (AL) network, the mushroom body (MB) and the lateral horn (LH). Plasticity is included at two levels, the AL and the MB. Plasticity in the interneurons of the AL allows for a decorrelation of odor response patterns within few trials, reducing generalization across odorants during training. Plasticity in the MB underlies the formation of associations. Uniglomerular projection neurons (PN) project from the AL to the MB calyx where each PN bouton forms the core of a microglomerular complex (Rybak and Menzel, 2010; Groh and Rössler, 2011) making contacts with several excitatory Kenyon cells (KC) as well and with inhibitory feedback neurons from the MB output. Based on physiological evidence (Sinakevitch et al., 2013) we assume reward-modulated plasticity in these inhibitory feedback synapses. The plasticity in this recurrent network structure leads to a training-induced modulation of bouton activity as observed in vivo (unpublished data, see previous chapter) and, in turn, modulates the pattern at the level of the MB output (Strube-Bloss et al., 2011). Subsequently, KC patterns are decoded by a MB-extrinsic neuron (EN) through a specialized dendritic tree that provides a layer of independent computational subunits (Polsky et al., 2004). As a result the MB output conveys a value code, i.e. odor identity represented in the combinatorial code of the PNs is recoded into odor valence at the MB output level (Strube-

Bloss et al., 2011; Menzel, 2014). A read-out of the value code by a downstream pre-motor neuron of the LH reflects the corresponding CR probability. We show that predictions derived from our network model match a range of behavioral data from elemental and non-elemental learning paradigms, including absolute and differential conditioning, trace conditioning, and odor patterning.

4.1 Introduction

The honeybee brain hosts a remarkable repertoire of associative learning faculties, and the link to cognition in higher animals has been made repeatedly (Giurfa and Giurfa, 2003; Menzel et al., 2006; Menzel, 2012). Learning under various experimental conditions has been characterized extensively by olfactory conditioning of the proboscis extension response (PER) (Bitterman et al., 1983; Giurfa and Sandoz, 2012; Matsumoto et al., 2012). In this classical conditioning paradigm, honeybees are subjected to odorants (conditioned stimuli, CS) that are presented in a defined temporal relation to a sucrose reward (unconditioned stimulus, US). The learning performance during training is monitored by computing the group-averaged probability for a proboscis extension (conditioned response, CR) on each conditioning trial.

The multitude of different variants of this protocol can be devided into absolute and non-absolute conditioning protocols. In absolute conditioning, honeybees are presented to only one conditioned stimulus during training. The effect of several experimental parameters on learning can be studied in this type of protocol, such as the concentration and duration of odorants (Pelz et al., 1997; Wright, 2004; Wright et al., 2005; Wright et al., 2009), the inter-stimulus-interval between odor and sucrose reward (Szyszka et al., 2011) or the inter-trial-interval (Menzel et al., 2001). Non-absolute conditioning on the other hand entails learning of stimulus-reward contingencies for more than one odorant during training. For example, in differential conditioning animals are faced with two different odors, only one of which appears in combination with a reward. This task can be made more difficult by reducing the concentration of odorants, or by employing a mixture of odorants at different ratios for the rewarded and unrewarded stimulus (Fernandez et al., 2009). Even more complicated variants of PER conditioning, such as negative or positive patterning (Deisig et al., 2001), introduce an ambiguity in the value of individual odor components by presenting them both in rewarded and unrewarded trials. Figure 4.1 summarizes the typical paradigm structure and the range of parametric variations included in this study.

In Drosophila, behavioral plasticity over consecutive training trials in individual animals is typically not monitored and considered further by theoretical studies, followind the assumption that the expression of behavior in a sample of fruit flies is homogeneous (Quinn et al., 1974). Consequently models are only constrained by group-averaged behavioral performance scores measured after training (Young et al., 2011; Wessnitzer et al., 2012). The case is different for classical conditioning in harnessed honeybees, where the experimental procedure allows to follow and record the behavior of each individual over training trials (Felsenberg et al., 2011; Matsumoto et al., 2012). These studies thereby provided the basis for studying inter-individual differences in learning performance in the honeybee (Scheiner et al., 2001; Scheiner et al., 2004; Pamir et al., 2014). Furthermore, knowledge on group heterogeneity supports the correct definition of behavioral model constraints (Pamir et al., 2011).

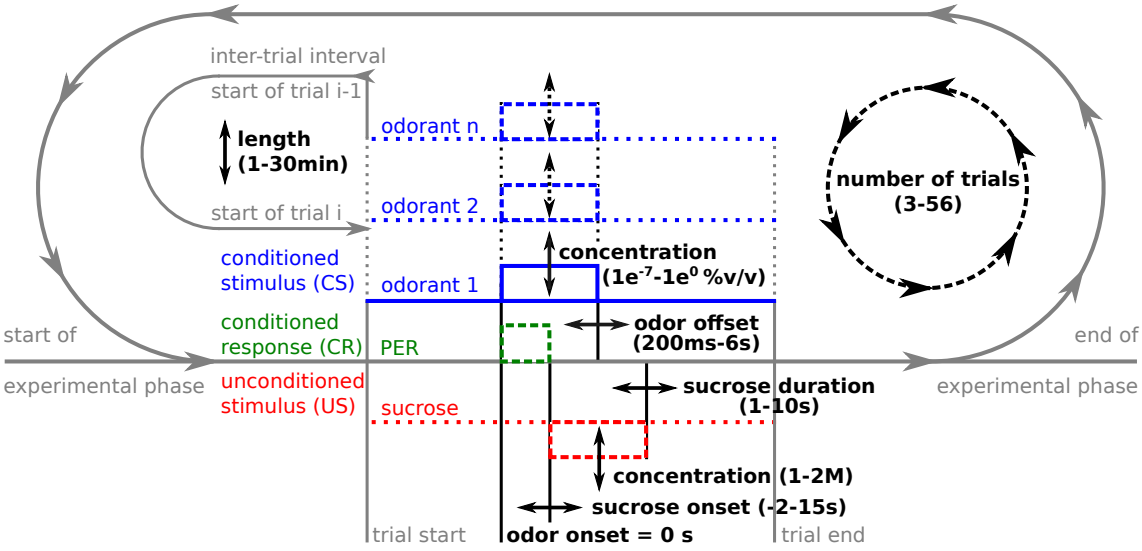

Figure 4.1: **General paradigm structure of olfactory conditioning of the PER.** Within an experimental phase, honeybees are typically subjected to a combination of one or more odorants (blue) for a given number of trials. These are combined with a sucrose reward (red) in a subset of the trials. During each trial the conditioned response (green) is monitored. Each experiment defines specific stimulus combinations with corresponding intensities and temporal relations. Values in brackets indicate the parametric ranges covered by conditioning experiments that were collected for this study.

The large body of evidence on behavioral plasticity in the honeybee is accompanied by numerous neuroanatomical, physiological, and biochemical studies on olfactory information processing and learning in the honeybee brain (Sandoz, 2011; Himmelreich and Grünewald, 2012, with a list of learning-related events in the honeybee brain in table 1). Despite the huge amount of data on trial-resolved behavioral plasticity on the one hand and neural information processing on the other, hardly any theoretical attempts have been made so far to explicitly link the two by a computational model, and those that exist focus on a specific experimental protocol (Bazhenov et al., 2013). Our study integrates several key findings on behavioral plasticity in the honeybee, which were observed under various different training protocols (Table 4.1). We explain the observed changes in conditioned response probabilities over consecutive training trials by neural computations in an abstract network model of the honeybee brain. To this end, we implemented an information processing scheme along the sensor-to-motor circuitry in the honeybee, which captures the current computationally relevant knowledge of honeybee physiology within an abstract network model. This approach allows us to investigate the effects of stimulus parameters, network structure, and computational principles on the emergence of behavior (Carandini, 2012).

Multiple existing modeling studies of the insect brain (Huerta and Nowotny, 2009; Young et al., 2011; Wessnitzer et al., 2012; Bazhenov et al., 2013) implement learning mechanisms between mushroom body intrinsic and mushroom body extrinsic neurons, a model structure that is

supported by physiological findings (Cassenaer and Laurent, 2012). In contrast to these, a series of physiological findings (Hammer, 1997; Szyszka et al., 2008; Haehnel and Menzel, 2010; Sinakevitch et al., 2013) suggests that an important mechanism for associative plasticity in the honeybee brain might be located in the microglomerular circuitry of the mushroom body calyx, mediated by inhibitory feedback neurons that integrate mushroom body intrinsic signals of the Kenyon cells and project this information back to the mushroom body calyx via the protocerebral tract (PCT, Fig. 4.2). We integrated this hypothesis into our network model complemented by a mechanism of neural plasticity located in the antennal lobe (AL), explained in detail in the following section. In section 4.3 we evaluate the learning performance of the system on the basis of a rich collection of data that characterizes behavioral plasticity in honeybees during olfactory conditioning. We found that predictions from our network model match a wide range of behavioral data from elemental and non-elemental learning paradigms. These results are discussed in the light of the implemented plasticity mechanism at inhibitory synapses of the MB calyx (section 4.4).

4.2 Methods

4.2.1 Network model

Figure 4.2 provides an overview of the most important neuron types involved in olfactory processing in the honeybee, most of which form the functional components of our network model. Odorant molecules interact with olfactory receptors (OR) in the antenna (ANT) and elicit distinct activity patterns across olfactory receptor neurons (ORN). These provide excitatory input to different types of neurons of the antennal lobe (AL), which is organized in spherically shaped neuropils called glomeruli. Two types of local neurons (LN), whose branching patterns stay within the AL and which are presumably mostly inhibitory, provide gain control through global inhibitory feedback (global LNs, gLN) as well as lateral inhibition through distinct asymmetrical connections (local LNs, lLN), respectively. Uniglomerular projection neurons (uPN) connect the AL with the lateral horn (LH) and the mushroom body (MB) via different antennal lobe tracts (ALT). Connections to the LH might be involved in the regulation of innate behavior and are not considered in our model. uPNs provide excitatory input to the MB-intrinsic Kenyon cells (KC). The activity patterns of the KCs are integrated by MB-extrinsic neurons (EN). While one subgroup of these forms an inhibitory feedback circuit back to the input region of the MB via the protocerebral tract (PCT, this abbreviation will also used for the corresponding group of neurons), other types of ENs send their axons to the LH. In addition, projections from multiglomerular projection neurons (mPN) of the AL also arrive at the LH. All of the three major neuropils that are involved in the processing of olfactory stimuli (AL, MB, and LH) are innervated by two octopaminergic neurons (ventral unpaired median neurons, VUM neuron) that mediate information on sugar reward. In the following paragraphs, we provide more information on neurophysiology and corresponding references. See Sandoz (2011) for a detailed review on olfactory perception and learning in honeybees.

Details of the network geometry are shown in figure 4.3. This includes numbers for each type of neuron, dimensions of connection matrices, and implemented sights of neural plasticity. Each processing step is explained in the subsequent sections. First, we introduce the Heaviside step function $\Theta()$ and the Kronecker delta function $\delta()$ that help to formalize the computational mechanisms we implemented:

$$\Theta(x) = \begin{cases} 1, x > 0 \\ 0, x \leq 0 \end{cases} \qquad \text{and} \qquad \delta(x) = \begin{cases} 1, x = 0 \\ 0, x \neq 0 \end{cases} \tag{4.1}$$

Two different versions of neural transfer functions are used, a normal sigmoid function $f()$ as well as a shifted and stretched sigmoid function $f_2()$. Both range from 0 to 1 and are parameterized for neuron group N by a threshold N_θ, a tuning broadness N_g and a center N_c

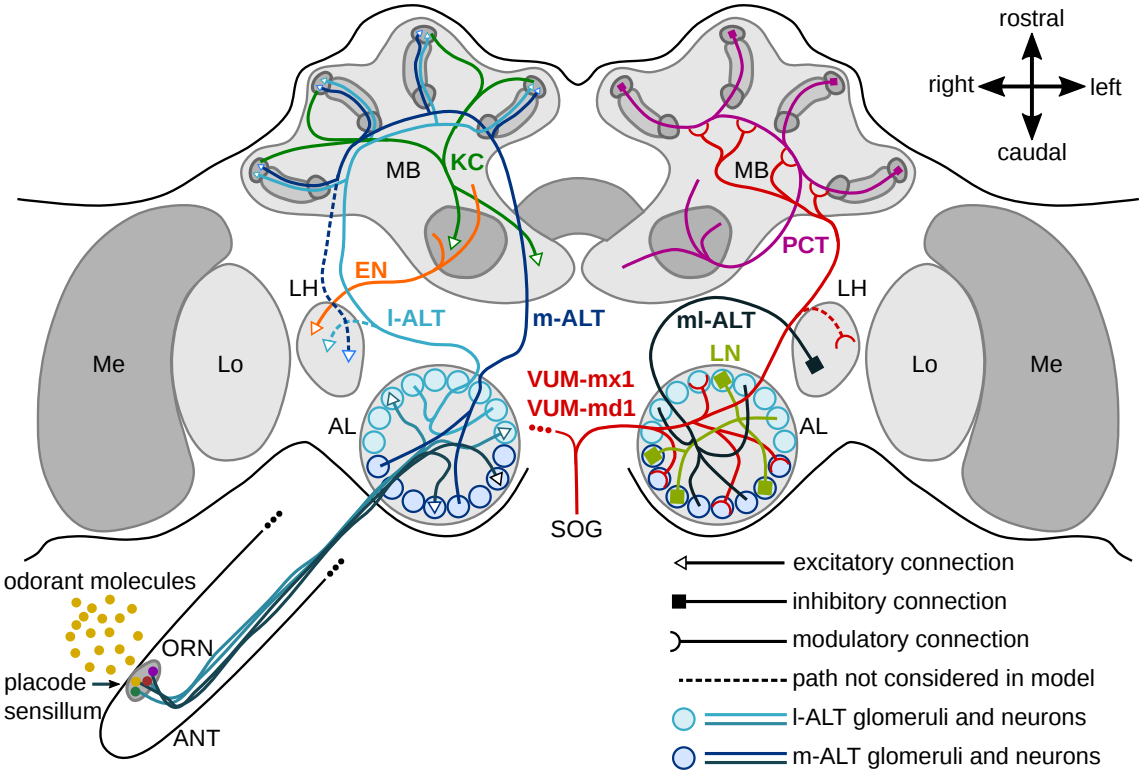

Figure 4.2: **Olfactory processing in the brain of the honeybee (after Sandoz 2011).** For clarity, neuron types are drawn in one hemisphere only. Receptors of olfactory receptor neurons (ORN) in the antenna (ANT) are stimulated by odorant molecules and excite neurons in the glomeruli (blue circles) of the antennal lobe (AL). Uniglomerular projection neurons project to the lateral horn (LH) and the mushroom body (MB) via two main pathways, the lateral and medial antennal lobe tract (l-ALT, m-ALT). Multiglomerular projection neurons provide inhibition to the LH via the medio-lateral ALT (ml-ALT). Lateral neurons (LN) are mainly inhibitory and stay within the AL. Two octopaminergic ventral unpaired median neurons (VUM) branch in AL, MB and LH. Kenyon cells (KC) receive excitatory input from uPNs. Their patterns are integrated by extrinsic neurons (EN) that send their axons to the LH. Next to these, feedback neurons integrate the KC pattern and project this information back to uPN-KC-connections in the MB calyx via the protocerebral tract (PCT).

(threshold and center default to 0 for some neuron groups):

$$f^N(x) = \frac{\Theta(x - N_\theta)}{1 + \exp(-N_g(x - N_c))} \tag{4.2}$$

$$f_2^N(x) = 2 \cdot \Theta(x - N_\theta) \cdot \left(\frac{1}{1 + \exp(-N_g(x - N_c))} - 0.5 \right) \tag{4.3}$$

Odor-reward integration

Each conditioning experiment consists of a series of trials (Fig. 4.1). For a given neuron group, a single activity pattern is calculated in each trial, neglecting any detailed temporal dynamics. The only exception to this is the temporal relation between the conditioned stimulus (CS)

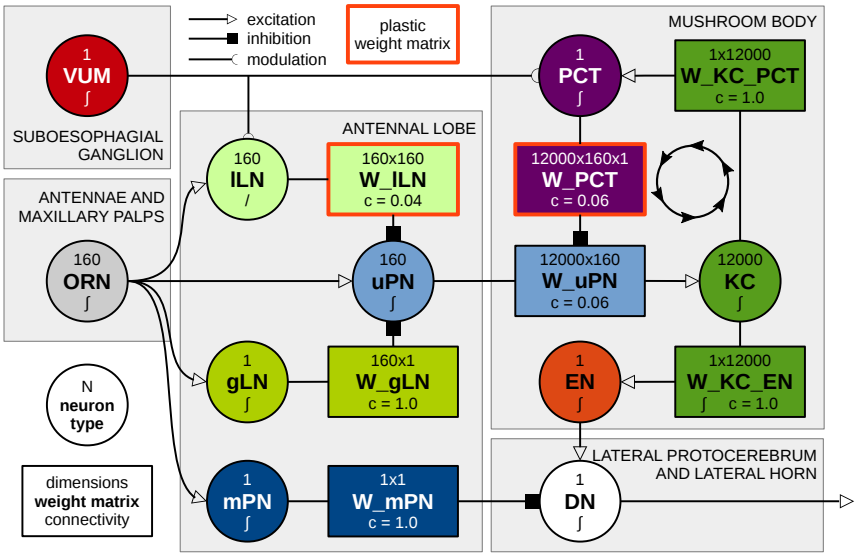

Figure 4.3: **Overview of the network structure.** Neuron groups are represented as circles including the number of computational unites per group, the group name and the type of transfer function (TF), with / for linear TF and $\int$ for sigmoidal TF. Rectangles represent weight matrices including dimensions, matrix name and connectivity (c). VUM: ventral unpaired median neuron, ORN: olfactory receptor neuron, lLN: local lateral neuron, gLN: global lateral neuron, mPN: multiglomerular projection neuron, uPN: uniglomerular projection neuron, KC: Kenyon cells, PCT: mushroom body feedback neurons following the protocerebral tract, EN: mushroom body extrinsic neurons, DN: downstream neuron.

and the unconditioned stimulus (US). Based on neurophysiological evidence (Hammer, 1993; Schröter et al., 2007; Sinakevitch et al., 2013), we assume this information to be present in the form of a synaptic tag Γ_{OA} at the two proposed sights of neural plasticity, i.e. synapses of lLNs in the AL and synapses of PCTs in the MB.

We calculate an eligibility trace $E(t)$ that mediates temporal information about the CS (Wessnitzer et al., 2012), independent of its concentration (Fig. 4.4). It increases exponentially after CS-onset and decreases exponentially after CS-offset, as characterized by the two variables α_E and τ_E:

$$E(t+\Delta t) = E(t) \cdot \exp^{-\frac{\Delta t}{\tau_E}} + (1 - E(t)) \cdot \alpha_E \cdot \Theta(CS(t)) \tag{4.4}$$

During a rewarded trial, the value of $E(t)$ at US-onset defines the magnitude of Γ_{OA}:

$$\Gamma_{OA} = \sum_{t_{start}}^{t_{end}} E(t) \cdot \delta(t - US_{onset}) \cdot \Delta t \tag{4.5}$$

We use this tag to scale the weight changes that occur as a consequence of the temporal stimulation pattern of each trial. We assume that these weight changes occur within the intertrial-interval (ITI, 1-30 min) and do not influence the tag of the subsequent trial ($E(t_0) = 0$).

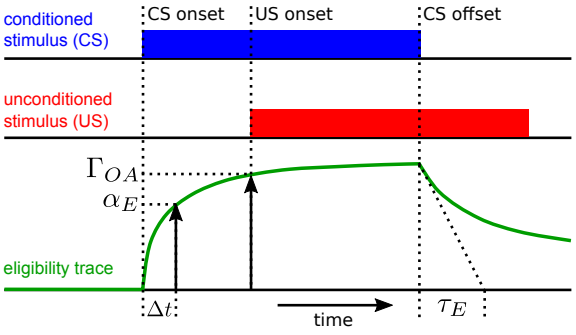

Figure 4.4: **Eligibitlity trace $E(t)$ and synaptic tag Γ_{OA}.** The eligibility trace increases exponentially during CS presentation and decreases exponentially after the CS ends. Speeds of increase and decrease are parameterized by α_E and τ_E, respectively. The temporal relation between CS and US determine the magnitude of the synaptic tag Γ_{OA}.

On the basis of the smallest time difference between CS- and US-onset that was used in the experiments we studied ($200ms$), we defined the temporal resolution to be $\Delta t = 0.1s$. α_E was determined according to three experimental groups from (Wright et al., 2009) that received the US 200, 500, and 800 ms after CS-onset, respectively. We fitted τ_E with respect to trace conditioning experiments reported in (Szyszka et al., 2011).

Antenna

In the model, the response of an odor receptor (OR) to a specific ligand is defined by its binding affinity k (the inverse of the receptor sensitivity), by its Hill coefficient m that determines the effective concentration range, and by the ligand concentration c as

$$f^{OR}(k,m) = \frac{1}{1+(\frac{k}{c})^m}. \tag{4.6}$$

Hence, k represents the center of a sigmoid function in relation to a concentration, and m controls the broadness of the corresponding sigmoid.

In a calcium imaging study, (Sachse and Galizia, 2003) measured odor response dynamics in the AL as a compound signal, which was presumably dominated by olfactory receptor neurons (ORN). They estimated both binding affinity and Hill coefficient for a number of glomeruli and odors. On the basis of this data, we generate the distribution of binding affinity k_i of receptor OR_i for a single ligand as

$$k_i = 10^{(k_{min}+k_{max}\cdot X_i^{OR_{exp}})}, X_i \sim U([0,1]) \tag{4.7}$$

With $OR_{exp} < 0$ this formalism accounts for the uneven distribution of receptor sensitivities with only few highly sensitive receptors. The Hill coefficient OR_m is shared by all olfactory

receptors in the model and is set to the average of the Hill coefficients reported in (Sachse and Galizia, 2003) (Figs. 4.6A, 4.12A). k_{min} determines the decimal logarithm of the minimal binding affinity (highest possible sensitivity), which we defined to be -5. The two remaining parameters were fitted with respect to the concentration dependence of the fraction of active ORNs (compound signal) estimated in (Sachse and Galizia, 2003) (Figs. 4.6D, 4.12B,D).

In each trial, the CS is defined by one or a mixture of odorants, i.e. ligands. While they always share their temporal properties, they might differ in their concentrations (Fig. 4.1). Indirect evidence for mixture perception of olfactory receptors in honeybees suggests hypoadditivity as the predominant mechanism (Deisig et al., 2006; Deisig et al., 2010). Assuming that hypoadditive and synergistic activities, which are evident in the AL of honeybees, are mainly caused by lateral interactions within the AL network, we implemented a hypoadditiv mixture integration for all receptors. Thus, in the model each olfactory receptor neuron ORN_i responds to a mixture of N_L ligands with the maximum of all ligand specific activations of its receptor k_i:

$$ORN_i = \max_{l=1}^{N_L}(k_i^l) \tag{4.8}$$

The overlap between activity patterns of two different odorants OR^i and OR^j is controlled by s_{ji} that defines the similarity between the patterns on a scale from 0 (opposite tuning) to 1 (identical tuning). The Nth tuning pattern OR^i is generated on the basis of all previously generated patterns using the same receptor activations shuffled in a new sequence $\mathcal{S}^i$:

$$\begin{aligned}\mathcal{S}^i = argsort\Bigg(\sum_{j=1}^{N-1}\Big(&\mathcal{S}_j^{=}\cdot[s_{ij}-0.5]^+ + \\ &\mathcal{S}_i^*\cdot\big(0.5-|\frac{1}{N-1}\cdot\sum_{k=1}^{N-1}s_{ik}-0.5|\big)+ \\ &\mathcal{S}_j^{\neq}\cdot[0.5-s_{ij}]^+\Big)\Bigg)\end{aligned} \tag{4.9}$$

$\mathcal{S}_j^{=}$ denotes the sequence of ligand affinities of odorant OR^j, $\mathcal{S}_j^{\neq}$ stands for the opposite sequence, and $\mathcal{S}_i^*$ is a uniformly randomly generated sequence. Therefore, each newly generated receptor activity pattern is a compromise between similarities to previously defined patterns (Fig. 4.5). With a total of 160 ORs and, in turn, 160 ORNs the model is approx. scaled to the number of potentially functional types of olfactory receptors and the number of AL glomeruli in honeybees (Sandoz, 2011). Table 4.3 contains a description of all ANT parameters.

Antennal lobe

The AL network is defined by four neuron groups, receiving the ORN activity pattern (Fig. 4.6A) as their input. gLNs and mPNs are both represented by a single nonlinearity that integrates

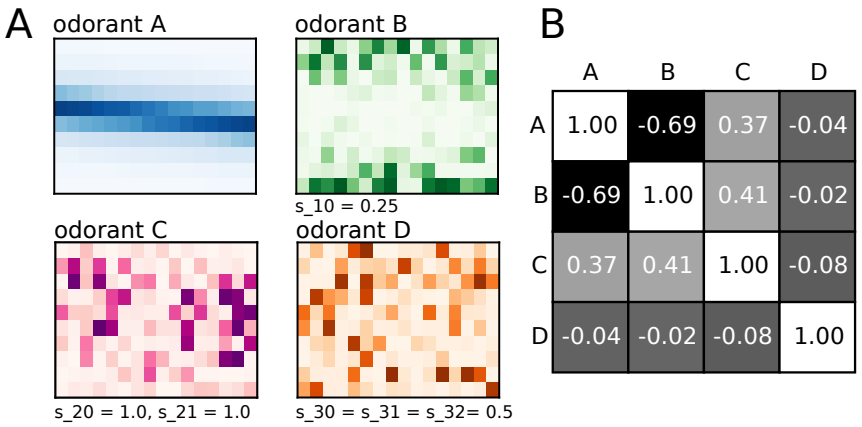

Figure 4.5: **Parametric configuration of overlap between odorants based on similarities s_{xy}.** A: Example of receptor tunings for multiple odorants. Tuning vectors of 160 receptors are shown as twodimensional matrices ordered with respect to the first odorant (blue). The second odorant (green) is anticorrelated to odorant A ($s_{10} < 0.5$). Tuning of the third odorant represents a compromise between the blue and green odorant. Tuning of the fourth odorant is not correlated to any others. Color intensity represents tuning strength. B: Coefficients of variation between tuning vectors corresponding to odorants in A.

the activity of all ORNs and provides gain control through inhibition for the AL and the LH, respectively (Fig. 4.6D, L). Lateral interactions between glomeruli are provided by lLNs (Fig. 4.6B). lLNs receive excitatory input from ORNs through a 1-to-1 connectivity and, effectively, copy the ORN activity pattern. lLNs connect to uPNs with a probability of $p^{lLN,uPN} = 0.25$. The majority (85%) of these connections is inhibitory and the corresponding weights are initially normally distributed around w_{in}^{lLN} and w_{ex}^{lLN} for excitatory and inhibitory connections, respectively (Fig. 4.6C). lLNs receive octopaminergic signals from the VUM neuron (Sinakevitch et al., 2013) that support the formation of Γ_{OA} and is assumed to be available to lLNs as an intracellular marker. Based on this, we implemented a mechanism of reward-dependent plasticity for these connections that is described in section 4.2.1.

uPNs of the two major tracts lALT and mALT are not distinguished in the model. A single computational unit represents the uPNs of each glomerulus (corresponding to the number of ORs, ORNs and lLNs). Average response latencies in LNs have been shown to be significantly shorter than in uPNs (Krofczik et al., 2008). Therefore, we calculate the corresponding activity pattern of uPNs on the basis of previously calculated LN activity:

$$uPN = f_2^{uPN}\left(uPN_{sr} + ORN + W^{lLN} \cdot lLN - gLN\right) \tag{4.10}$$

uPN_{sr} is a spontaneous rate that we added based on observations for PNs in honeybees (Deisig et al., 2003; Nawrot et al., 2010; Meyer et al., 2013). In the model, its function is to provide a baseline activity that facilitates the emergence of activity patterns in the MB at low odor concentrations. The parameters of the uPN transfer function were determined regarding data

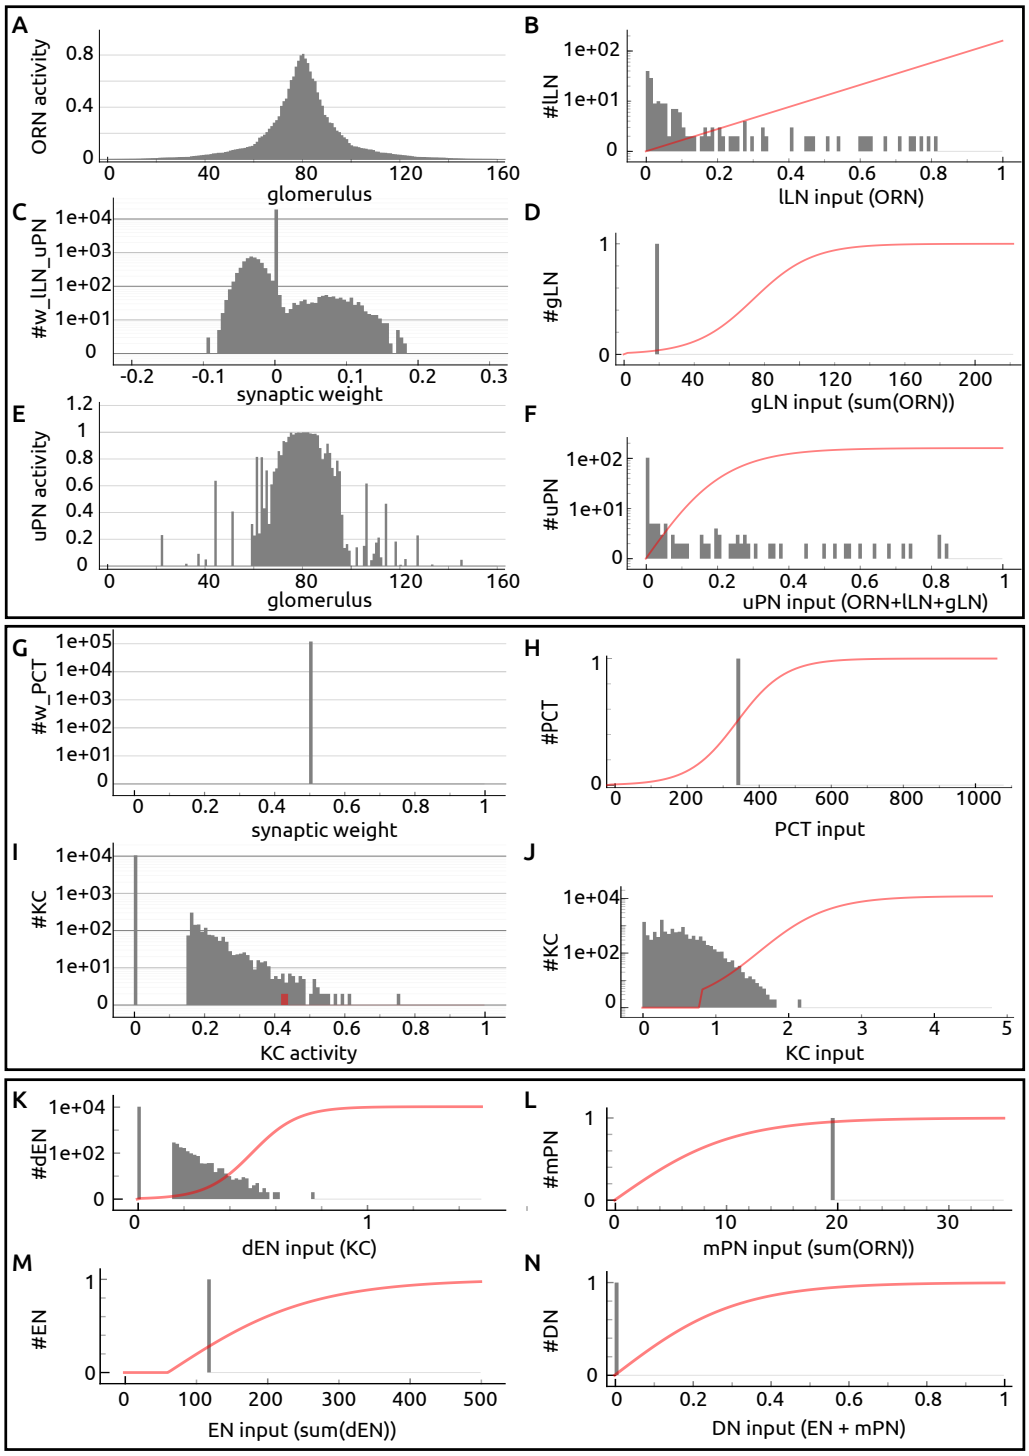

Figure 4.6: **Different stages of neural processing in the network model.** Red curves represent neural transfer functions and are always scaled to unity. A-F: Antennal lobe. G-J: Mushroom body. K-N: Lateral horn. For more informations please refer to the corresponding sections.

from fruit flies (Bhandawat et al., 2007) (figures 4.15). Subsequently, we defined values for the concentration dependence of the fraction of active uPNs (Sachse and Galizia, 2003). This allowed us to fit the concentration dependent level of gain control in the antennal lobe by adjusting the three parameters gLN_g, gLN_c, and w_{gLN} (Fig. 4.14). The resulting low activity in gLN (Fig. 4.6D) matches the observation that application of GABA in the AL of honeybees silences all uPN activity (Sachse et al., 2002). Table 4.3 contains a description of all AL parameters.

Mushroom body

The model contains 12000 MB-intrinsic KCs that receive excitatory input from uPNs. The corresponding connectivity matrix is parametrized by $N_{PN,KC}$ that defines the average number of KCs a single uPN projects to, resulting in a set of 120,000 synaptic terminals of uPNs that contact the population of KCs. These are uniformly randomly distributed across all possible synaptic connections. To quantify population sparseness in KCs we used the measure of kurtosis (Field, 1994).

KC activity patterns are integrated by two neuron groups, PCTs and ENs. Both are implemented as a single nonlinearity (Fig. 4.6H, M). The EN connects the MB with the LH and is further described in section 4.2.1. The PCT projects back to the MB calyx and contacts each synaptic connection between uPNs and KCs through an inhibitory synapse, forming an inhibitory feedback loop:

$$KC = f^{KC}(W^{uPN} \cdot uPN - W^{PCT} \cdot PCT) \tag{4.11}$$

$$PCT = f^{PCT}(W^{KC,PCT} \cdot KC) \tag{4.12}$$

$$= f^{PCT}\Big(W^{KC,PCT} \cdot f^{KC}(W^{uPN} \cdot uPN - W^{PCT} \cdot PCT)\Big) \tag{4.13}$$

These equations can be solved numerically and result in a stable solution that represents a point of balance between KC and PCT activity. The weights of $W^{KC,PCT}$ were scaled to 1, the plastic weights of W^{PCT} were always initialized with 0.5 (Fig. 4.6, B).

PCTs are also contacted by the octopaminergic VUM neuron, as suggested by physiological evidence (Sinakevitch et al., 2013). We assume that this connection provides relevant information about the temporal relation between CS and US and, in turn, supports the formation of Γ_{OA} that is available to PCTs as an intracellular marker (see section 4.2.1). Based on this, we implemented a mechanism of reward-dependent plasticity for these connections that is described in section 4.2.1.

Some of the parameters of the MB network were set to heuristically chosen values. In addition, we defined target values for the concentration dependent PCT activity and fraction of active KCs and fitted the remaining parameters according to these (Fig. 4.16). Table 4.3 contains a

description of all MB parameters.

Lateral horn

A pool of ENs connects the MB with the LH. Analogous to PCTs, these are abstracted as a single computational unit with a nonlinear transfer function that integrates the activity of all KCs (Fig. 4.6M). However, the KC signals are individually transformed by a dendritic tree that provides a layer of independent computational subunits (Fig. 4.6K). These sigmoidally modulate the KC inputs prior to a global summation, a model that was proposed in (Polsky et al., 2004) to explain data recorded from cortical pyramidal neurons in vertebrates.

Yet another nonlinear unit, the DN, combines the EN signal with the activity of inhibitory mPNs (Fig. 4.6N). Here, mPNs connect the AL with the LH (Abel et al., 2001; Sinakevitch et al., 2013) and provide a global signal of AL activity (Fonta et al., 1993). With its DN activity the network generates a signal that represents the probability of a conditioned response (p(CR)).

Plasticity in the mushroom body

We apply a reward-dependent learning rule to the inhibitory PCT synapses that modulate the connections between uPNs and KCs on the basis of Γ_{OA} as well as pre- and postsynaptic signals. According to equations 4.14-4.18, the synaptic weight $w_{i,j}$ that modulates the connection between uPN_i and KC_j is decreased in the presence of a synaptic tag ($\Gamma_{OA} > 0$) and increased otherwise ($\Gamma_{OA} = 0$). The speed of this process is scaled by learning rates $\alpha^+_{w^{PCT}}$ and $\alpha^-_{w^{PCT}}$. In addition, following multiplicative rules w^{PCT}_{max} and w^{PCT}_{min} provide upper and lower boundaries for the synaptic weights, respectively. This mechanism leads to a disinhibition of rewarded patterns and a suppression of unrewarded patterns.

$$\Delta w^-_{i,j} = \Gamma_{OA} \cdot uPN_i \cdot KC_j \tag{4.14}$$

$$\Delta w^+_{i,j} = (1 - \Theta(\Gamma_{OA})) \cdot uPN_i \cdot KC_j \tag{4.15}$$

$$\Delta w^{PCT}_k = \alpha^+_{w^{PCT}} \cdot \Delta w^+_{i,j} \cdot (w^{PCT}_{max} - w^{PCT}_k) - \tag{4.16}$$

$$\alpha^-_{w^{PCT}} \cdot \Delta w^-_{i,j} \cdot (w^{PCT}_k - w^{PCT}_{min}) \tag{4.17}$$

$$w^{PCT}_k(t + \Delta t) = w^{PCT}_k(t) + \Delta w^{PCT}_k \tag{4.18}$$

$\alpha^-_{w^{PCT}}$ was fitted on the basis of a simple absolute conditioning protocol (Pamir et al., 2011) to match the generally observed rapid learning dynamics in honeybees during classical conditioning (Pamir et al., 2014). $\alpha^+_{w^{PCT}}$ was fitted to the rate of extinction reported in (Stollhoff et al., 2005). Consequently, the same synapses that are reduced in efficacy during appetitive learning will be strengthened during unrewarded trials.

Plasticity in the antennal lobe

We integrated a reward-dependent learning rule into the inhibitory synapses of lLNs that modulate the activity of uPNs in the AL on the basis of Γ_{OA} as well as pre- and postsynaptic signals. In contrast to the plasticity model of the MB, pre- and postsynaptic signals are scaled by the difference between the expected reward -expressed by activity of DN - and the received reward. The synaptic weight w_k^{lLN} of neuron lLN_k that receives input from ORN_i and inhibits uPN_j is changed according to the following equations:

$$\mathcal{H}_k = ORN_i \cdot uPN_j \cdot |DN - \Gamma_{OA}| \tag{4.19}$$

$$\Gamma_{OA} > 0,\, \Gamma_{OA} - DN > 0.05 : \Delta w_k^{lLN} = \begin{cases} -w_k^{lLN} \cdot \alpha_{w^{lLN}} \cdot \mathcal{H}_k & \text{for} \quad w_k^{lLN} < 0 \\ (w_{ex,max}^{lLN} - w_k^{lLN}) \cdot \alpha_{w^{lLN}} \cdot \mathcal{H}_k & \text{for} \quad w_k^{lLN} > 0 \end{cases} \tag{4.20}$$

$$\Gamma_{OA} = 0,\, DN > 0.05 : \Delta w_k^{lLN} = \begin{cases} -(w_{in,max}^{lLN} + w_k^{lLN}) \cdot \alpha_{w^{lLN}} \cdot \mathcal{H}_k & \text{for} \quad w_k^{lLN} < 0 \\ -w_k^{lLN} \cdot \alpha_{w^{lLN}} \cdot \mathcal{H}_k & \text{for} \quad w_k^{lLN} > 0 \end{cases} \tag{4.21}$$

$$w_k^{lLN}(t + \Delta t) = w_k^{lLN}(t) + \Delta w_k^{lLN} \tag{4.22}$$

Thus, in dependence on Γ_{OA} and DN, excitatory and inhibitory connections are changed in opposite directions. Positive weights are increased while negative weights are decreased for rewarded patterns, and vice versa for unrewarded patterns. Most of the involved parameters were determined with respect to behavioral data from a differential conditioning experiment (Pamir et al., 2011).

4.2.2 Behavioral model constraints

Classical conditioning of the proboscis extension response in the honeybee was introduced more than fifty years ago (Takeda, 1961; Giurfa and Sandoz, 2012). Details on this experimental procedure can be found elsewhere (Scheiner et al., 2001; Felsenberg et al., 2011; Matsumoto et al., 2012). In order to parameterize and evaluate our computational model, we collected a large set of classical conditioning protocols from the experimental literature (Table 4.1). Our study concentrates on the dynamics of behavioral plasticity that appears over the course of a few conditioning trials, roughly within up to a few hours. Hence, a prerequisite for any protocol to be used was the availability of corresponding data.

Typically, during a standard absolute conditioning experiment, a distinct subgroup of 10-20% of the bees does not react to the rewarded odor in any of the conditioning trials (Pamir et al.,

2011). When possible, we excluded these non-responding animals, which results in a more homogeneous sample of animals with respect to learning performance. This procedure was not applicable to datasets in which the learning task was made more difficult by reducing the CS-concentration and/or the US-duration (Wright et al., 2009) or by increasing the inter-stimulus interval (Szyszka et al., 2011). In these datasets we subtracted a hypothetical basic level of non-responders, which was estimated based on those groups of bees that were part of the same experiment - and therefore were subjected to the same general conditions - but were not faced with an increased degree of difficulty.

Another subgroup of bees is characterized by the fact that they extend their proboscis as a reaction to the CS presentation during the first conditioning trial. Recently it has been shown that these spontaneous responders share the same learning dynamics as bees that start responding in later trials (Pamir et al., 2014). However, in several datasets these spontaneous responders were excluded by the experimenter. Additionally, we did not explicitly integrate a mechanism of spontaneous response into our model, and therefore removed these individuals from other datasets as well if possible.

In the majority of the protocols the experimental raw data at the level of individual animals was available to us. In these cases, we were able to exclude non-learners and spontaneous responders from the data as described. For some conditioning protocols the raw data was not available. To obtain targets for the simulations we transcribed the CR probabilities manually from the original publications. In total, we collected the data for 89 experimental groups from 17 studies corresponding to a total number of 3944 animals (Table 4.1). In some cases we combined experimental groups within a study that only differed with respect to long-term effects tested after several hours, for these are not considered here.

Parameter search

We fitted the parameters of the model in a number of steps to the available physiological data, to a subset of the behavioral data, and where necessary to heuristically defined targets. To this end, we implemented a simple brute search parameter optimization algorithm that minimizes the root mean squared error (RMSE) between the observed and the simulated conditioned response probabilities. For model m and a given set of data d from a classical conditioning protocol lasting N_T trials this is defined as

$$RMSE_{d,m} = \sqrt{\frac{1}{N_T} \cdot \sum_{T=1}^{N_T} p_d(CR(T)) - p_m(CR(T))}. \tag{4.23}$$

In some cases, the parameter search was be performed by minimizing the RMSE for a collection D of N_D classical conditioning protocols, scaled by the relative number of animals n_d in each

Reference	Amount of included groups	Amount of included animals	Raw data	Used to fit parameters	Description of learning paradigm
ELEMENTAL LEARNING PARADIGMS					
(Pamir et al., 2011)	1	61 (100)	√	√	absolute
(Marter et al., 2014)	1 (4)	158 (221)	√	–	absolute, US duration
(Wright et al., 2009)	21	877 (987)	√	√[1]	absolute, concentration
(Szyszka et al., 2011)	1	62 (93)	√	–	absolute, delay
(Szyszka et al., 2011)	7	263 (328)	√	√	absolute, trace
(Stollhoff et al., 2005)	1 (5)	177 (228)	√	√	absolute, extinction
(Chandra et al., 2010)	7	115 (124)	√	–	latent inhitibion
(Pamir et al., 2011)	1	87 (120)	√	√	differential
Yamagata, unpublished	1	13 (18)	√	–	differential, extinction
(Smith, 1998)	3	90	–	–	absolute, feature positive discrimination
NON-ELEMENTAL LEARNING PARADIGMS					
(Fernandez et al., 2009)	10	319	–	–	differential, concentration
Szyszka, unpublished	1	57 (89)	√	–	negative patterning
(Komischke et al., 2003)	2	64	–	–	positive/negative patterning
(Deisig et al., 2001)	6	422 (422)	√	–	positive/negative patterning
(Deisig et al., 2002)	6	245 (255)	√	–	positive/negative patterning
(Deisig et al., 2003)	6	237 (271)	√	–	positive/negative patterning,
(Chandra and Smith, 1998)	9	215	–	–	differential, neg. patterning, biconditional discrimination
Total	82 (89)	3462 (3944)	13/17	6/17	

Table 4.1: **Overview of classical conditioning studies.** For groups, numbers in brackets indicate the amount of experimental groups in the original study that we combined to a single group. For animals, numbers in brackets indicate the number of bees that were part of the corresponding study before we corrected for non-learners and spontaneous responders. [1]six groups were used for parameter fitting.

data set d:

$$RMSE_{D,m} = \sum_{d=1}^{N_D} RMSE_{d,m} \cdot n_d \tag{4.24}$$

During parameter searches for the concentration dependency of both uPNs and KCs, heuristic target values were introduced in equation 4.24 as individual error terms (Figs. 4.14 and 4.16). To summarize, model parameters were fitted as follows:

- We fitted ANT, AL and MB parameters to match concentration dependent activities of neural groups in each stage (ORNs, uPN and KCs, respectively). Target values were based in part on physiological data.

- Using a learning rate of $\alpha_{PCT} = 1.0$, we fitted LH parameters with regard to the observed p(CR) of three absolute conditioning groups from (Wright et al., 2009) that differed with respect to CS concentration (0.3, 0.003 and 0.00003 %v/v).
- We fitted the temporal dynamics of the synaptic tag Γ_{OA} to match the observed p(CR) in three absolute conditioning groups from (Wright et al., 2009) (varying CS duration to fit α_E) and all trace conditioning groups from (Szyszka et al., 2011) (varying US onset to fit τ_E).

After we determined this first set of parameters, we refined the learning dynamics of the model in two steps:

- Based on the hypothesis that plasticity in the MB calyx is the predominant mechanism underlying rapid learning dynamics in honeybees, we refined the learning rates that control MB plasticity ($\alpha^{+}_{w^{PCT}}$ and $\alpha^{-}_{w^{PCT}}$) to match acquisition and extinction dynamics during absolute conditioning (Pamir et al., 2011; Stollhoff et al., 2005). This step involved the deactivation of AL plasticity. We refer to this configuration where only MB plasticity is active as m_0.
- Based on the hypothesis that plasticity in the AL is the predominant mechanism underlying odor decorrelation, we refined the parameters that control dynamics of AL plasticity ($w^{lLN}_{ex,max}$, $w^{lLN}_{in,max}$, and $\alpha_{w^{lLN}}$) to match data from differential conditioning (Pamir et al., 2011). Both plasticity mechanisms are active in this case. We refer to the corresponding configuration as m_1.

Model evaluation

For both configurations, we generated model predictions for all collected conditioning protocols. Here, the same set of parameters was used for both configurations, differing only in terms of AL plasticity being switched off (m_0) or switched on (m_1). Some protocols used mixtures of odorants. For these, we adjusted the parameters that control the odor tuning to match the observed odorant similarity as evident from generalization effects during the first block of trials. For those components of the network model that were generated using a standard pseudo-random number generator (receptor tuning, connectivity matrices), we used predefined seeds for each of these components to initialize the number generator. Thus, we used the identical network instance for both model configurations.

4.3 Results

We implemented a computational model of odor processing and learning in the honeybee that comprises a mechanism of synaptic plasticity located in the calyx of the mushroom body, a network scheme that has so far not been studied in a computational framework (Fig. 4.3). We also collected a large assembly of behavioral data from classical conditioning protocols (Table 4.1). In this section we try to establish a quantitative link between the neural computation in the model and the observed dynamics of behavioral plasticity.

In order to quantify the difference between model predictions and honeybee data we calculated the root mean squared error (RMSE) of the difference in the probability of a conditioned response $p(CR)$ across learning trials for each protocol (section 4.2.2). For brevity, we present detailed results only for a subset of protocols in the following sections. Details about the corresponding conditioning protocols can be found in table 4.4. Table 4.2 lists the average $RMSE$ of all considered studies to provide a comprehensive overview.

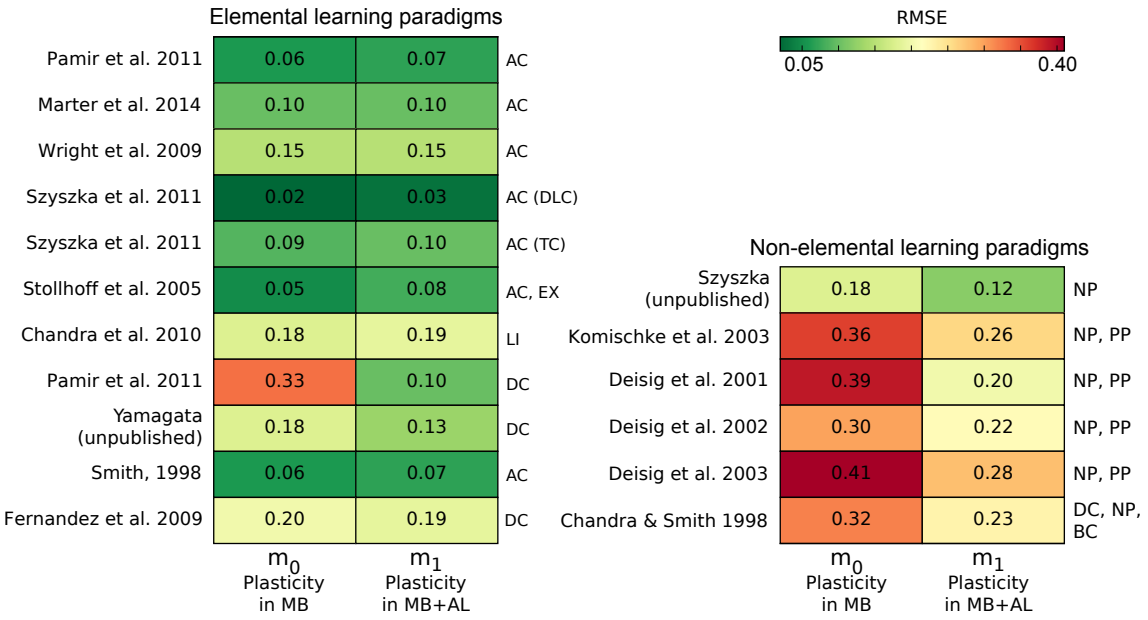

Table 4.2: **Average model performances.** Model performances were characterized by the root mean squared error (RMSE) calculated for the difference in probabilities of conditioned responses between simulation results and bee data, averaged across all experimental groups from a given study. AC: absolute conditioning. DLC: delayed conditioning. TC: trace conditioning. EX: extinction. LI: latent inhibition. DC: differential conditioning. NP: negative patterning. PP: positive patterning. BC: biconditional discrimination.

4.3.1 Results of elemental learning paradigms

Elemental learning paradigms are characterized by the fact that a discernible property of a rewarded stimulus will never appear in an unrewarded trial. Hence, a simple link between this property and the reward suffices to solve the task (Giurfa, 2003).

Absolute conditioning

The typical variability between identically treated groups of bees during absolute conditioning calculated as RMSE amounts to approx. 0.05-0.1, as estimated on the basis of the experimental groups from Stollhoff et al., 2005 and Marter et al., 2014 (average group size ~40). Within the margin of this variability, the learning dynamics during acquisition of nearly all absolute conditioning protocols were reproduced by the model. The differences between the two model configurations m_0 and m_1 are negligible for absolute conditioning.

The dependence on concentration and CS-duration was fitted according to a subset of groups from (Wright et al., 2009). Here, the RMSE is higher than in the other absolute conditioning experiments, a result that can be attributed to a combination between the way we determined parameters for the neural groups in the lateral horn (LH) and a rather high variability in the corresponding behavioral data (Fig. 4.7A-F).

In latent inhibition, a stimulus is presented during a pre-training phase in the absence of reward for a given amount of trials, followed by a number of rewarded training trials (Chandra et al., 2010). As evident from the corresponding behavioral data, honeybees respond with a delay in acquisition that is dependent on the number of unrewarded pre-training trials. This effect is reproduced by the model for up to 20 pre-training trials (Fig. 4.7G, H). For groups of bees that were subjected to even more pre-training trials, acquisition speed is further reduced and, judged on the basis of the 6th training trial, also the asymptotic level of the conditioned response is significantly reduced (Fig. 4.7H). This effect is not captured by the model, causing an average RMSE near 0.2.

The time constant of the decay in eligibility τ_E was determined on the basis of data from a set of trace conditioning protocols, in which an increase in the temporal gap between CS-offset and US-onset led to a reduced acquisition speed (Szyszka et al., 2011). Therefore, it is not surprising that this salient effect is paralleled by the trial-averaged $p(CR)$ of the model (Fig. 4.7J-L). But further more, also the linear increase in acquisition is closely reproduced by the model.

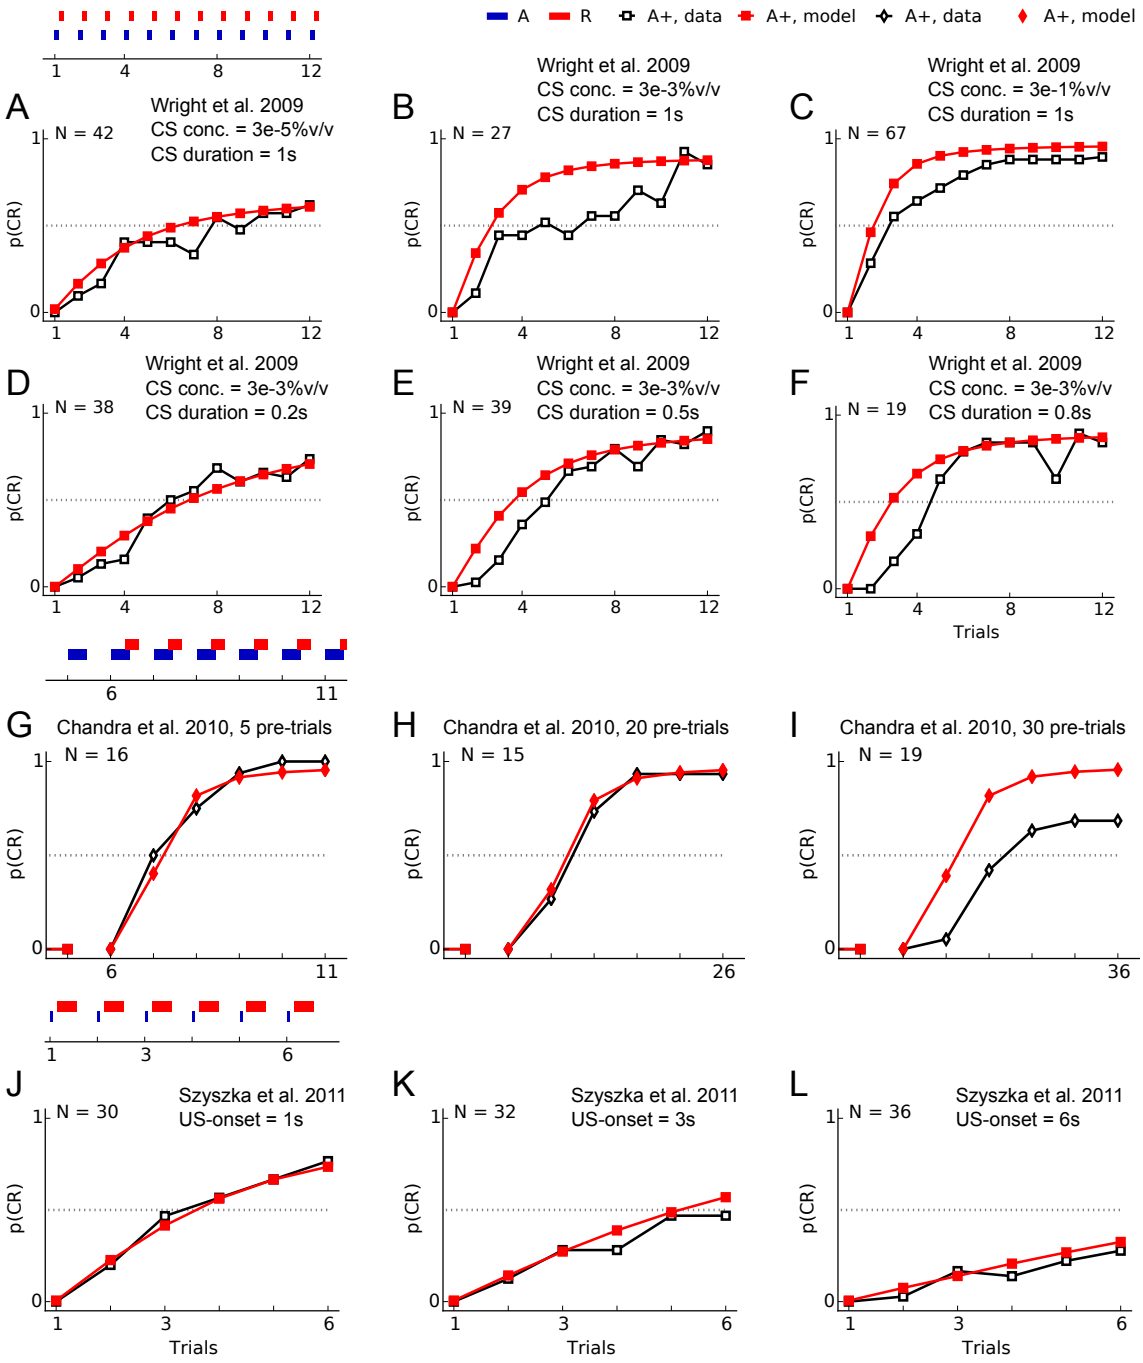

Figure 4.7: **Observed and simulated dynamics of associative learning in selected protocols of absolute conditioning**. Black curves show the experimentally observed CR probabilities. Red curves show the simulated CR probabilities of model configuration m_0. Stimulus onsets and durations in each trial are illustrated by colored pictograms. CS-duration varied among A-C, which is not shown in the corresponding pictogram.

Plasticity in the mushroom body during absolute conditioning

Figure 4.8 presents details about the MB network during absolute conditioning. In the model, MB plasticity is expressed by a decrease of efficacy in a subset of PCT synapses (section 4.2.1). This process is illustrated by a decrease of the average synaptic weight of W^{PCT} (Fig. 4.8A). As a consequence, two opposing effects exist in the mushroom body network. While the activity of some KCs is drastically increased (Fig. 4.8F), the total number of active KCs goes down from more than 13% and levels out at below 12% after three trials (Fig. 4.8E). This is caused by an increase in the overall level of KC activity (Fig. 4.8C) and, in turn, an increase in the overall level of inhibition at the MB input (Fig. 4.8B). Concerning the population sparseness of the KC pattern, the model predicts a strong increase during the first three training trials followed by a constant decrease during later trials (Fig. 4.8D).

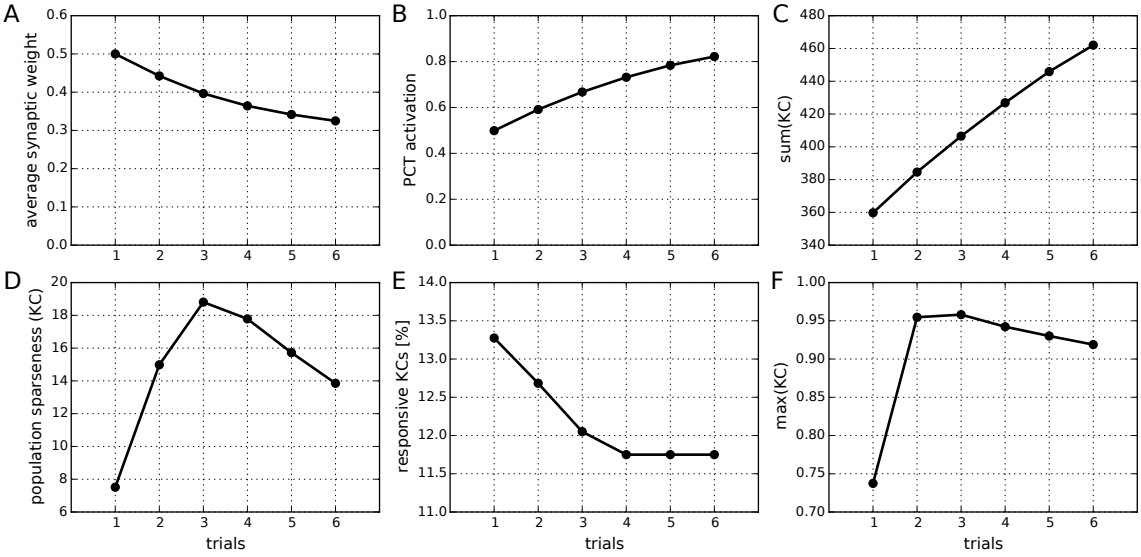

Figure 4.8: **Effects of plasticity in the MB during absolute conditioning.** A: Synaptic weights of connectivity matrix W^{PCT} are initialized at 0.5. A subset of weights is reduced during learning. Average values were calculated for these weights. B: Inhibitory feedback neurons respond with an increased activity during learning, caused by an increase in overall KC activity shown as the sum across all KCs in C. D: The model predicts two phases of change in population sparseness of KCs, a rapid increase during the first trials followed by a slower decrease thereafter. E: The ratio of responsive KCs is reduced during learning and saturates at just below 12%. F: The activity of a subset of KCs is strongly increased during learning. Through an increase in inhibition at the MB input, the maximum activity in KCs is constantly slightly reduced during the last three trials.

Differential conditioning

In differential conditioning, one rewarded odor and one unrewarded odor are presented in turn. Typically, the rewarded odor is presented in the first trial, followed by the unrewarded odor in the second trial. Depending on the perceptual similarity of the used odors, a number of

bees will generalize and respond with a CR in the second trial. Between the two datasets from (Pamir et al., 2011) and (Yamagata, unpublished) the ratio of bees that generalized differs significantly (Fig. 4.9A-D). For each of these protocols, we chose the similarity between the two tuning vectors (parameterized by s_{10}) to match the observed generalization in the second trial. In both cases, plasticity in the MB alone is not sufficient to separate the rewarded from the unrewarded odor as effectively as observed in the behavioral data (Fig. 4.9A, C), especially if the similarity between tuning vectors is high (Fig. 4.9A). Using the model configuration m_1 that comprises the additional AL plasticity, the model traces more closely resemble those of the observed data (Fig. 4.9B, D).

The quantitative measure of model performance that we used (see section 4.2.2) can be misleading in some cases, as evident from the learning paradigm studied in (Fernandez et al., 2009). It involves the presentation of two odorants in different combinations of concentration ratios. Further more, the CS-duration was varied across some of the experimental groups. Individual data was not at hand and, consequently, we could not control for spontaneous responders that were part of some groups. Interestingly, in this study all bees received the same pseudorandomized sequence of rewarded and unrewarded stimuli. Here, we present the behavioral data and model results for two exemplary groups (Fig. 4.9E-H).

In the first example from this study, the ratio of odorant concentrations was 10:0, i.e. the two used odorants did not appear as a mixture. Within each trial the CS was presented for 4 s with an overlap of 1 s between CS and US presentation (Fig. 4.9F, H). About 15% of the bees spontaneously responded in the first trial. The high level of generalization towards the unrewarded odor suggests a high perceptual similarity between the used odors. It reaches an even higher level in subsequent trials. For the rewarded odor the response level saturates after about four presentations (trial 7), due to the rather low odor concentration used (approx. 0.003 %v/v). In both configurations, the model response to the CS+ also saturates after trial 7, although at a slightly higher level. The effect of odor decorrelation through plasticity in the AL (m_1) is visible in the last two unrewarded trials (Fig. 4.9H).

In the second example, bees had to learn to distinguish between odorants that were mixed in a concentration ratio of 9:1, with one of the two possible combinations of concentration levels being assign to CS+ and the other one to CS- (Fig. 4.9E, G). In addition, odors were only presented for 500 ms with a gap of 500 ms between CS-offset and US-onset. At the end of this rather demanding task, a low level of differentiation appears in the behavioral data. Focusing on the rewarded odor, the model prediction of m_0 (no AL-plasticity) does not reach the same levels of p(CR) but closely follows the general shape of the behavioral data. For the unrewarded odor the difference in response levels is larger but, more importantly, model responses stay above those of the rewarded pattern. This shows that the model is not able to separate both stimuli from one another. In the case of m_1 (distributed plasticity) the general shape of the responses

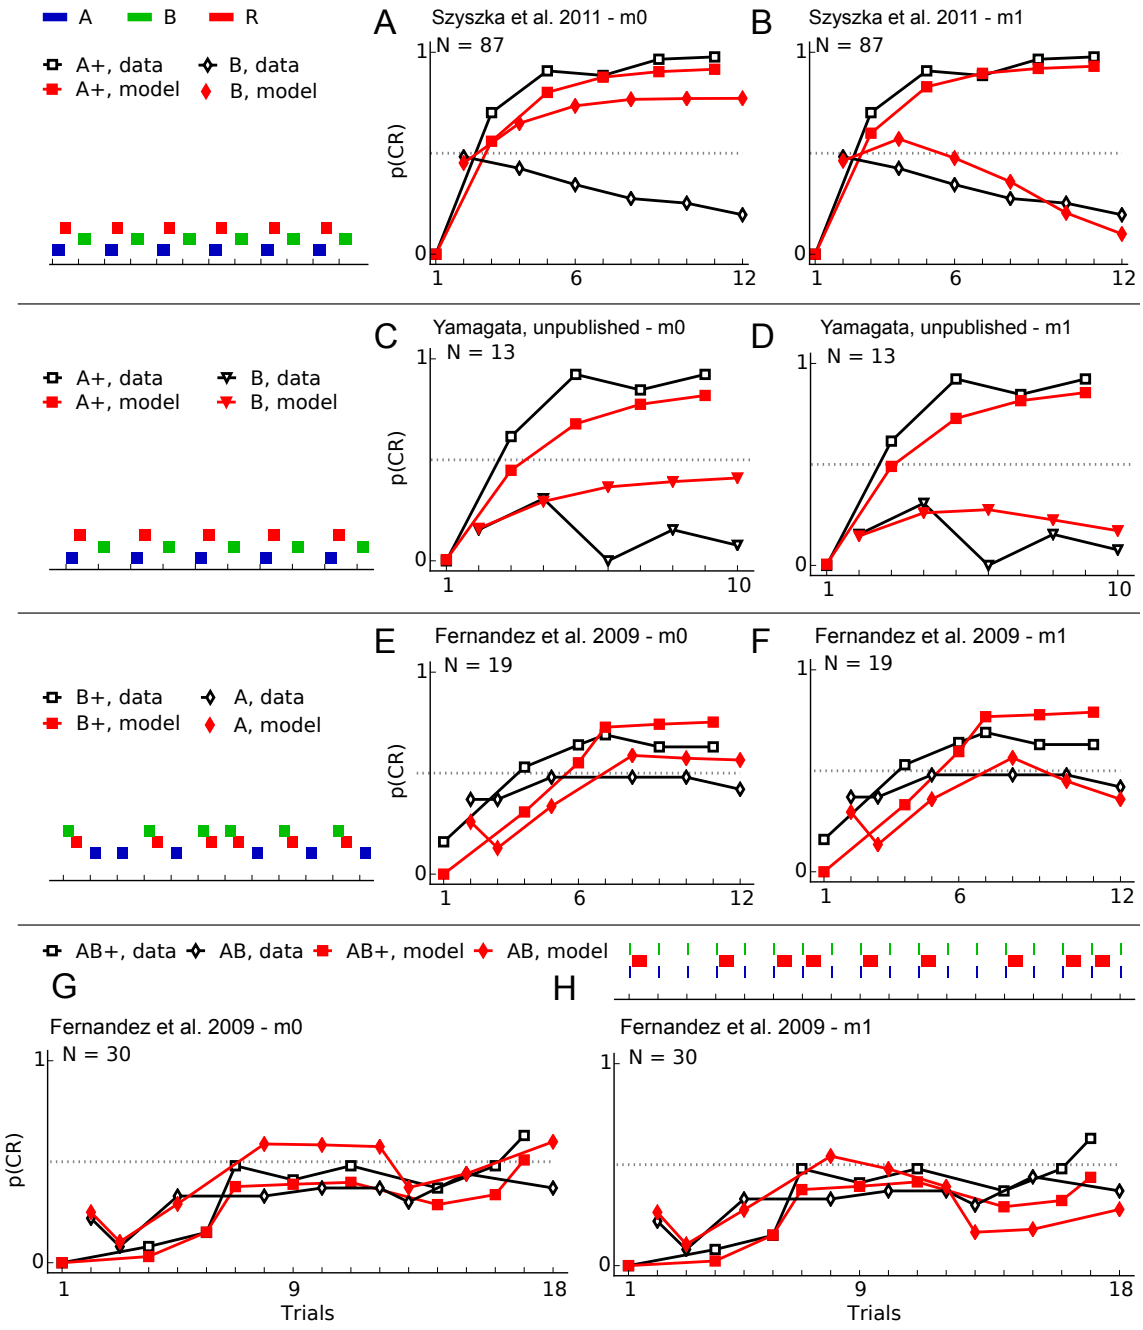

Figure 4.9: **Observed and simulated dynamics of CR probabilities in selected protocols of differential conditioning**. Black curves show the experimentally observed CR probabilities. Red curves show the simulated CR probabilities of model configuration m_0 (A, C, E, and G) and m_1 (B, D, F, H). Stimulus onsets and durations in each trial are illustrated by colored pictograms.

is similar to m_0, but after trial 12 the response levels for the rewarded odor stay above those of the unrewarded odor. Although the difference in response levels at the end of the experiment is not quite as strong as in the behavioral data, the qualitative outcome is the same. This qualitative difference in model performance is not expressed in the quantitative measure we used.

Plasticity in the antennal lobe during differential conditioning

Figure 4.10 illustrates details of the AL network during the differential learning paradigm shown in figure 4.9A,B. The modeled mechanism of AL plasticity involves a tuning of the weight matrix W^{lLN} that modulates the uPN activity pattern (section 4.2.1). Briefly, for rewarded patterns excitatory connections are strengthened while inhibitory connections are depressed. Conversely, unrewarded activity in the AL leads to a decrease in excitatory weights and an increase in inhibitory weights. For the given protocol these opposing processes are shown in Figs. 4.10A and B. Over time, the weight changes become smaller and the synaptic weights level out at higher absolute values for both positive and negative weights. The effect of this tuning is shown in Figs. 4.10C and D. A general increase in inhibition underlies the reduction in the average uPN activity pattern for both odors (Fig. 4.10C). However, the pattern of the rewarded odor A+ is only slightly suppressed. In comparison to this the pattern of the unrewarded odor receives much more inhibition. Effectively, the initially rather high correlation between both odor patterns ($C_V \approx 0.6$) is clearly reduced over time ($C_V \approx 0.4$).

4.3.2 Results of non-elemental learning paradigms

Learning paradigms are classified as non-elemental when stimuli ambiguously appear in both rewarded and non-rewarded training trials (Giurfa, 2003). Though complexity of training protocols can vary substantially, the number of training trials is typically larger than in elemental learning paradigms. In most of the paradigms pseudorandomized trials sequences were used that differed between subgroups of bees from the same conditioning protocol. In these cases we could not infer the exact sequences that were used on the basis of the available data. To compensate for this, we calculated the average $p(CR)$ of the model across multiple simulations of pseudorandomized trial sequences.

A clear difference in model performance between m_0 and m_1 is visible in this collection of learning paradigms (Table 4.2). While m_0 could not solve most of the more difficult tasks of negative and positive patterning, the qualitative outcome of the observed behavior was reproduced by m_1 in all cases. Still, the RMSE between model prediction and behavioral data reached high values for m_1 in some non-elemental learning paradigms (Komischke et al., 2003; Deisig et al., 2003). This can be mainly attributed to a combination of the large number of trials and the degree of odor decorrelation in negative patterning that resulted at times in a faster and stronger separation of rewarded and unrewarded stimuli in the model prediction.

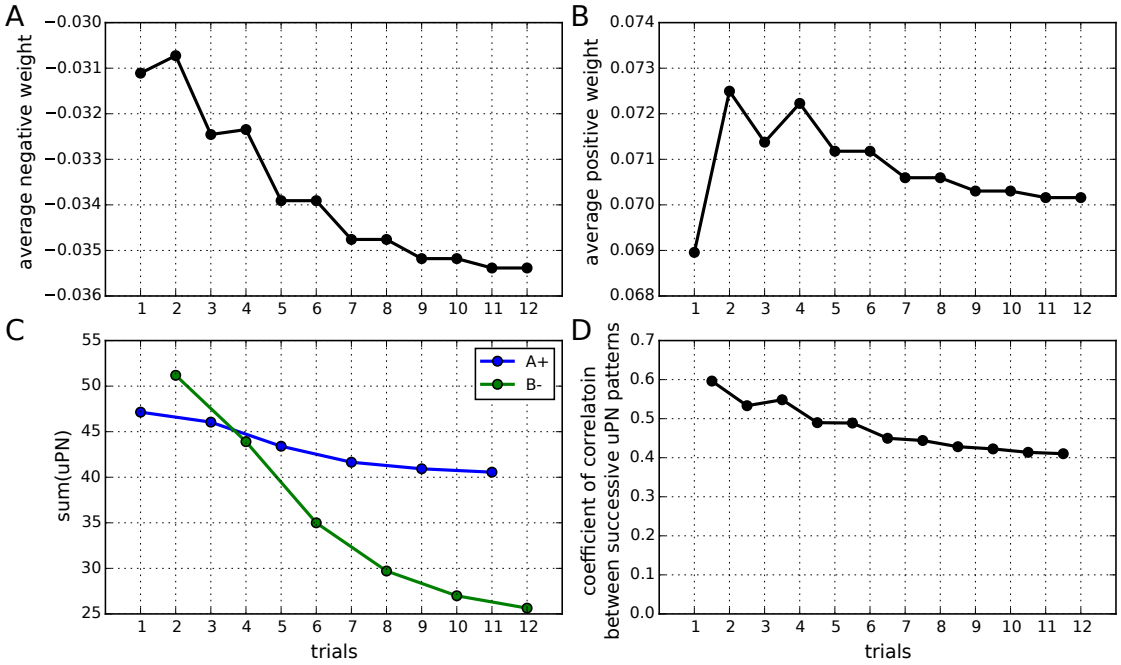

Figure 4.10: **Effects of plasticity in the AL network during differential conditioning.** A: The strength of inhibitory weights (absolute value) is decreased after rewarded trials (uneven numbered trials) and increased after unrewarded trials (even numbered trials). On average, inhibitory weights are increased, i.e. acquire higher negative values. B: The strength of excitatory weights is increased after rewarded trials and decreased after unrewarded trials. C: Rewarded odorant A is only slightly suppressed by the increased inhibition in the AL, while the average uPN activity towards unrewarded odorant B is drastically reduced. D: The correlation between uPN activity patterns is reduced through plasticity in the AL during differential learning.

Positive patterning

We selected a typical example for positive patterning that was studied in (Deisig et al., 2001). In this learning protocol, the mixture of two odorants was paired with a US (AB+) while presentations of the individual components were not (A-, B-). Each of the three stimuli appeared equally often in a sequence of 24 pseudorandomized trials (Fig. 4.11A,B). Beginning with the second presentation of the rewarded odor, there is a small but clear separation in honeybee response levels between the rewarded pattern and the unrewarded components. This separation gradually increases during subsequent trials, caused by a gradual decrease in response probabilities for the unrewarded stimuli, while the response levels for the CS+ remain around 0.5. Using only MB plasticity, model configuration m_0 cannot disentangle the unrewarded stimuli from the rewarded one. In the end, it responds to the single odorants as strongly as to the pattern, far away from the behavioral data. In contrast to this, with AL-plasticity switched on (m_1) the model reaches nearly the same levels of response for all of the stimuli, though it takes slightly longer to suppress CS-. Qualitatively, this result is exemplary for other learning paradigms of positive patterning.

Negative patterning

The example of negative patterning also involved two odorants. These were rewarded in trials where only one of them was presented. Thus, the odor pattern was not paired with a reward. In a pseudorandomized trial sequence, unrewarded stimuli appeared twice as often as rewarded (Fig. 4.11C, D). After the first block of six trials, an obvious separation between rewarded and unrewarded stimuli exists in the behavioral responses. For model configuration m_0, the MB plasticity mechanism is sufficient to separate rewarded from unrewarded stimuli, but this separation is less distinct (Fig. 4.11C). With the additional odor decorrelation in the AL, the suppression of the unrewarded odorant pattern follows a similar path as in the behavioral data, but the response levels for the rewarded odorants even surpass those observed in the bees. In general, this outcome is paralleled by other learning paradigms of negative patterning.

Biconditional discrimination

Figures 4.11E and F contain an example of biconditional discrimination studied in (Chandra and Smith, 1998). Four odorants were presented in four different pairwise combinations. Only two of these combinations were paired with sugar as a positive reinforcement. Specific to this study, the experimenters applied a salt solution to the antenna as punishment in trials that were not to be rewarded, a mechanism that we did not explicitly integrate into the network model. Over the course of 24 trials, each odorant appeared as often in rewarded as in punished stimuli. Bees learned to distinguish between rewarded and punished stimuli and acquired a strong positive association towards both rewarded pairs of odorants. A clear separation of rewarded and unrewarded stimuli is visible in the results of model configuration m_0, but response levels for unrewarded patterns were far higher than observed in bees (Fig. 4.11E). In combination

with AL plasticity though, response levels of m_1 quickly dropped in unrewarded trials to values even lower than those observed in bees.

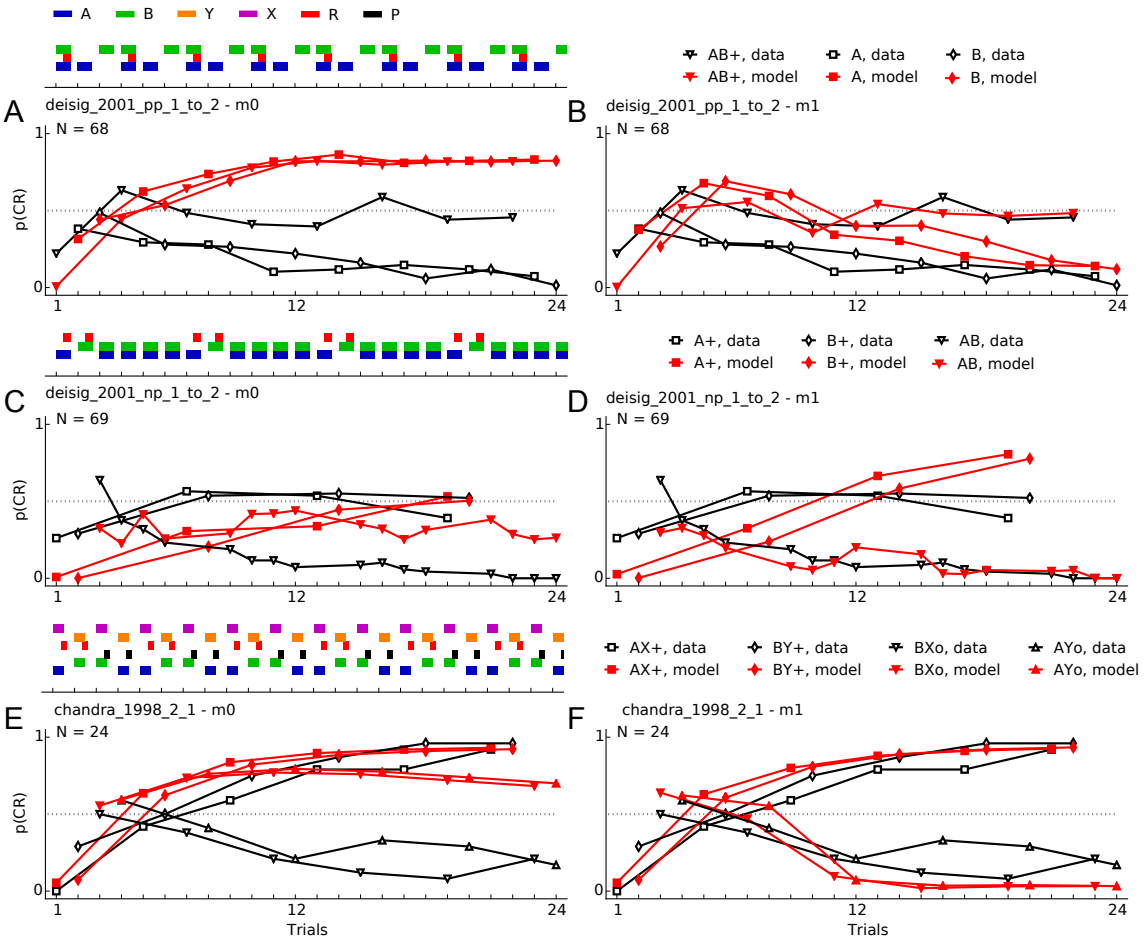

Figure 4.11: **Observed and simulated dynamics of CR probabilities in selected protocols of non-elemental conditioning paradigms.** Black curves show the experimentally observed CR probabilities. Red curves show the simulated CR probabilities of model configuration m_0 (A, C, and E) and m_1 (B, D, and F). Stimulus onsets and durations in each trial are illustrated by colored pictograms.

4.4 Discussion

We investigated olfactory processing and learning characterized by data from a range of classical conditioning of the proboscis extension response (PER) in honeybees with a coherent computational network model. Individual processing stages of our model were designed to capture the essential details of the physiological properties of the bee brain. We integrated a reward dependent plasticity mechanism as inspired by physiological findings (Sinakevitch et al., 2013) that entails the modulation of inhibitory feedback in the mushroom body (MB) calyx. This hypothesis has been formulated before (Szyszka et al., 2008; Haehnel and Menzel, 2010) but has so far not been implemented within a computational model. We generated predictions on honeybee behavior for a large collection of both elemental and non-elemental learning paradigms and compared these with the original data to evaluate our model.

We found that this form of plasticity in the MB (model configuration m_0), if calibrated on the basis of learning dynamics observed in absolute conditioning and extinction trials, is compatible with a large number of elemental learning paradigms but fails to effectively separate odorants in differential an more challenging non-elemental conditioning protocols. However, in combination with a plasticity mechanism that supports decorrelation of odorants in the antennal lobe (AL, configuration m_1) the model successfully reproduces the qualitative outcome in all elemental and non-elemental conditioning paradigms.

This result does not mean that an effective decorrelation mechanism could not also be computed by the proposed MB circuitry: Increasing the "unlearning" rate for unrewarded patterns $\alpha^{+}_{w^{PCT}}$ in an arbitrarily strong suppression. As a consequence though, the learning rate $\alpha^{-}_{w^{PCT}}$ that supports disinhibition of rewarded patterns needs to be increased to balance both mechanisms. Our result does rather support the hypothesis of a distributed plasticity across AL and MB without ruling out theories of locally restrained foci of plasticity, which needs to be investigated further.

In regard to our quantitative measure of model performance it should be noted that our approach suffers from the high degree of variability across different PER conditioning experiments (Matsumoto et al., 2012; Frost et al., 2012). More so, between identically treated groups of the same study we estimated the average variability to be similar to the average variability in prediction error during absolute conditioning, even when excluding known sources of heterogeneity such as non-learners and spontaneous responders (Pamir et al., 2011). Considering that we were not able to correct for these in many of the non-elemental learning paradigms, an increased prediction error is to be expected here. The existing differences between model prediction and behavior could be significantly reduced by increasing the number of degrees of freedom in the model for individual studies or even individual experimental groups. But considering the complexity of the model, such an approach could have likely resulted in some degree of overfitting of the behavioral data and would have hindered the assessment of the general validity of the model.

The process of finding the best general model of associative learning during PER conditioning in honeybees naturally requires a complete set of behavioral data. Several important training conditions are missing in the current study, such as blocking (Smith and Cobey, 1994; Gerber and Ullrich, 1999), reversal learning (Hadar and Menzel, 2010), backward conditioning (Hellstern et al., 1998), conditioned inhibition, or second-order conditioning (Takeda, 1961). Although we are continuously trying to enlarge our collection of behavioral data, for many of the older studies data acquisition of learning dynamics often proofs to be impossible, especially with respect to individual data.

Given the collection of learning paradigms at hand, many of these include training and test trials covering a wide range of time scales, which we did not give credit to in our model. We configured our model to represent learning dynamics that appear within a period of up to a few hours. Thus, plasticity in our model is primarily scaled to represent learning in the range of mid-term memory. Behavioral performance is, though, controlled by multiple processes of memory formation and decay (Hammer and Menzel, 1995). Therefore, the predictive power of the model will drop for paradigms that resolve different forms of memory in the behavioral data. An example for this phenomenon might have occurred in the study on latent inhibition from (Chandra et al., 2010). As described in the results, model predictions closely matched behavioral data for a number of up to 20 unrewarded pre-trials. Groups of bees that received even more pre-trials (30, 40, 50) showed a significantly decreased acquisition speed and possibly a reduced asymptote in CR probabilities. This effect was not captured by our model. Considering the corresponding durations of three and more hours, this divergence from model predictions might indicate the existence of a type of long-term memory that governed the observed behavioral plasticity in part.

Assuming the existence of the proposed mid-term memory trace in inhibitory synapses of the MB calyx, our model makes predictions about the activity of the involved neuron types during appetitive learning that might be tested experimentally. As presented in the results (Fig. 4.8), the model predicts an overall increase in population sparseness across KCs. This study revealed two mechanisms underlying this process: The activity of a specific subset of KCs is drastically increased. At the same time, other KCs that were originally responsive towards the rewarded odor become silent. Thus, a high appetitive value is not necessarily encoded by high PCT activity, but rather entails a strongly increased activity of a subset of KCs.

A key assumption of the proposed plasticity mechanism for the decoding of the KC pattern at the MB peduncle - the output region of the MB - is the non-linear integration by individual dendritic branches and the subsequent integration of the filtered KC pattern that is provided by the MB extrinsic neuron (EN). Thus, the EN provides a much more selective filter for high KC activities than a linear summation would (Polsky et al., 2004). Such a nonlinear dendritic

transformation at the level of individual synapses could be provided by the single Pe1 neuron that exists in each hemisphere of the bee brain and arborizes throughout all layers of the MB peduncle (Rybak and Menzel, 1998). Furthermore, the Pe1 receives input from PCT neurons and has repeatedly been found to decrease its activity on account of associative learning on timescales compatible with the debated processes (Okada et al., 2007; Hussaini and Menzel, 2013). This form of reward dependent plasticity in the Pe1 could result from a varying ratio between excitatory KC activity and inhibitory PCT activity as observed during absolute conditioning in our model (Fig. 4.8B, C). However, considering the lack of detailed knowledge about connectivity in the lateral horn (LH) and in the name of minimizing the number of unknown parameters we omitted the explicit integration of the Pe1 into our model at this stage.

The level of detail in terms of both network structure and temporal resolution could be readily increased. However, the model already allows to investigate theories of neural coding during compound processing or to estimate the potential of novel conditioning protocols. Apart from that, it offers multiple options to improve its current predictive power. To name one, the dynamics of mixture separation in m_1 during non-elemental paradigms could be adjusted to the learning dynamics observed in behavior. This would entail handling the most intriguing aspect of the model: The interdependency of the two distributed plasticity mechanisms remains to be analyzed.

4.5 Supplemental material

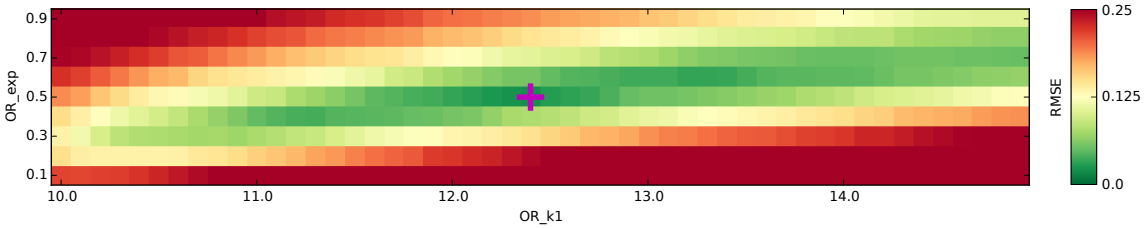

Figure 4.12: **Parameter search for ANT parameters.** Both parameters control the distribution of binding affinities across all 160 olfactory receptors. OR_{k1} defines the upper boundary of binding affinities. OR_{exp} controls the random placement of binding affinities between lower and upper boundary (OR_{k0} and OR_{k1}). A value near 1 results in uniformly distributed binding affinities. A value lower 1 reduces the probability of low binding affinities, i.e. high receptor sensitivities.

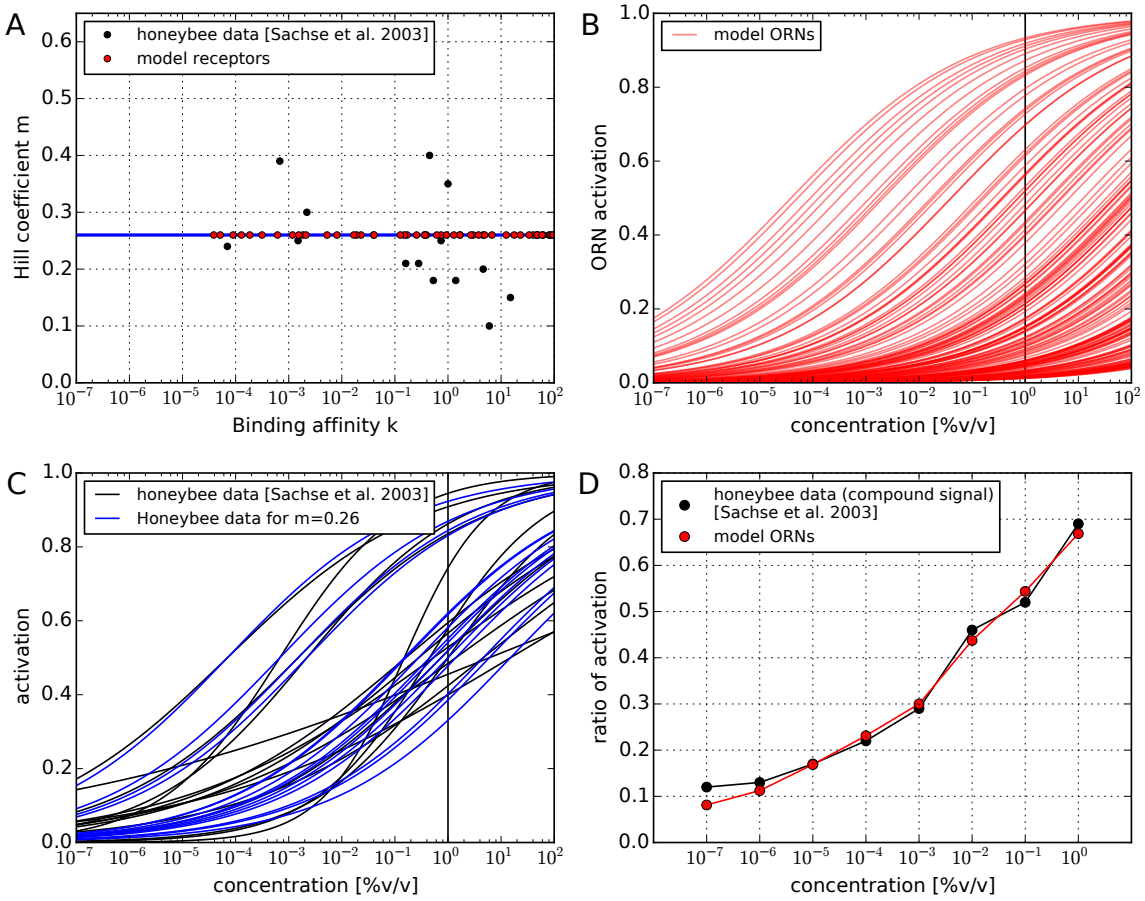

Figure 4.13: **Binding affinities of olfactory receptors (OR) and resulting concentration dependence of olfactory receptor neurons (ORN).** A: Blue line marks the average Hill coefficient of receptor dynamics reported in (Sachse and Galizia, 2003). Model receptors were randomly placed along this line (red dots). B: Concentration dependent activations of all 160 model ORNs resulting from the placement ORs shown in A. C: Comparison between concentration dependencies of ORNs as reported in (Sachse and Galizia, 2003) (black curves) and as resulting from a constant Hill coefficient (blue curves). D: Concentration dependent ratio of active ORNs in the model (red dots) and as reported in (Sachse and Galizia, 2003) (black dots).

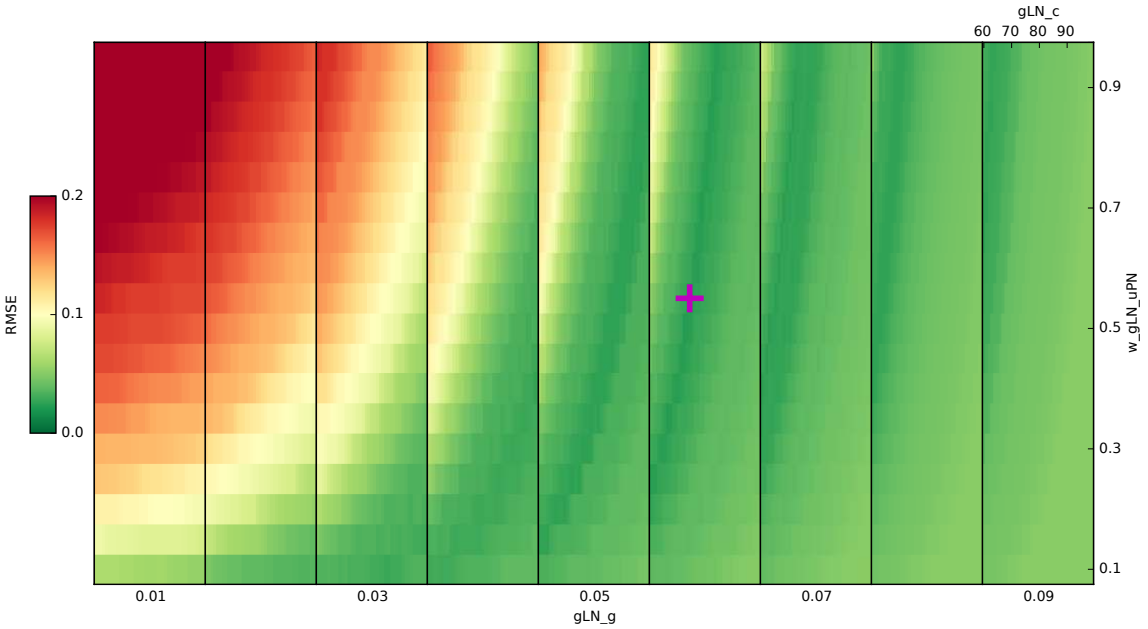

Figure 4.14: **Parameter search for AL parameters.** After fixing all other parameters, gain control through gLNs was fitted on the basis of target values for the concentration dependent ratio of uPNs. This included the synaptic weight $w_{gLN,uPN}$ (identical for all synapses) and the shape of the lLN transfer function, configured by the center gLN_c and broadness gLN_g of a sigmoid.

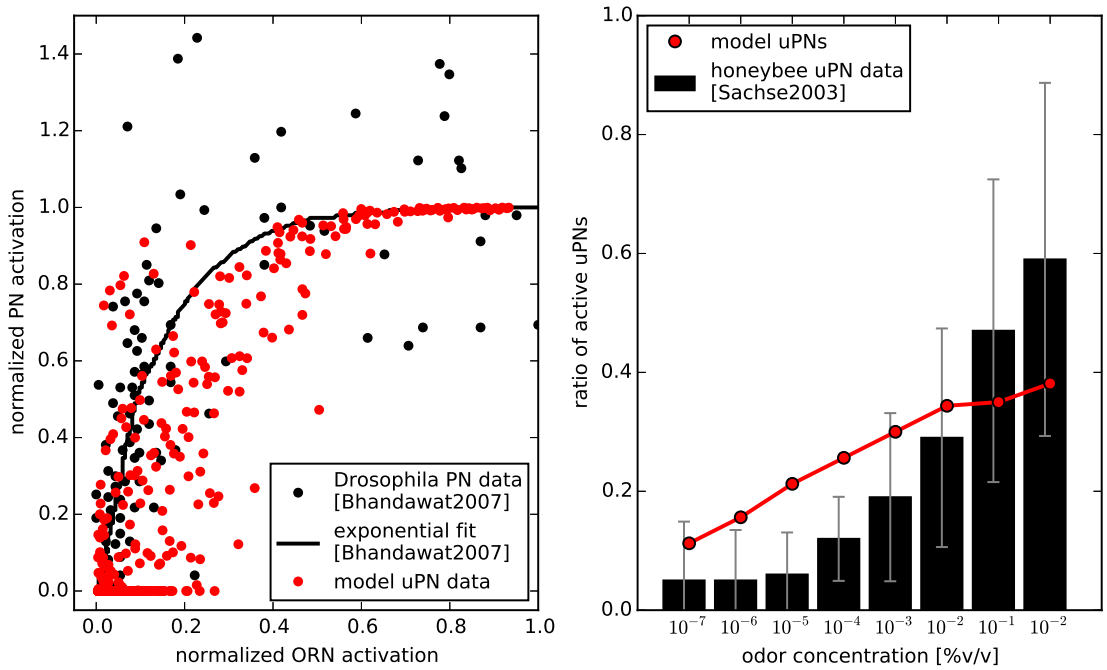

Figure 4.15: **Transfer function and concentration dependent activation of uPNs.** Left: Transfer function of uPNs was defined with respect to PN data from drosophila. Right: Concentration dependence of the ratio of active uPNs. Gray lines represent standard deviation of bee data.

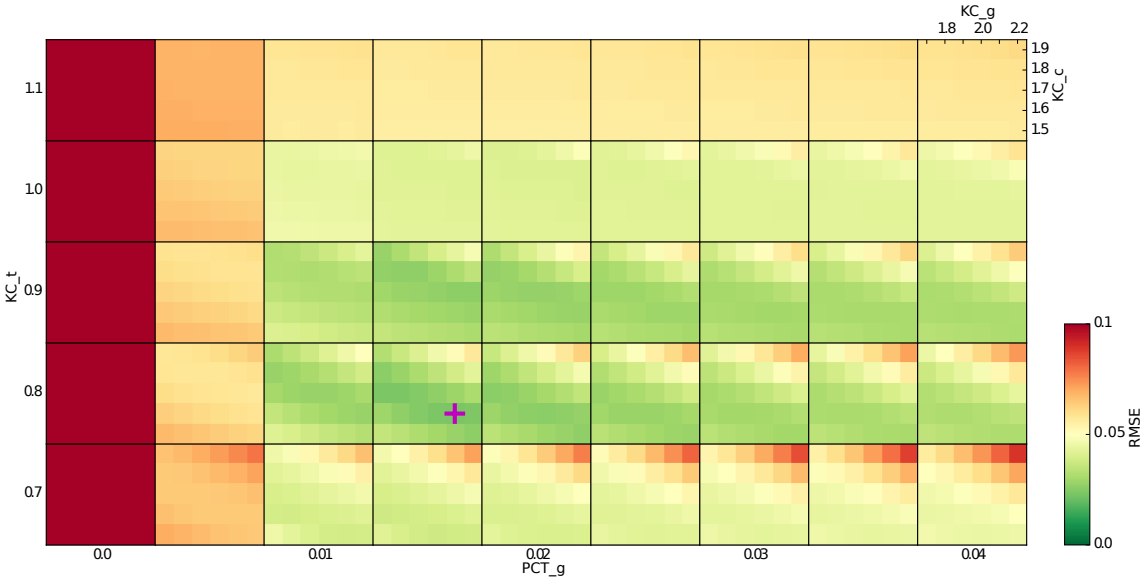

Figure 4.16: **Parameter search in the MB.** First, PCT_g was heuristically defined. Subsequently, the remaining four parameters that shape the transfer function of KCs (KC_c. KC_g, KC_t) and PCT (PCT_g) were fitted through a parameters search. Target values were heuristically defined ratios of active KCs (4-14%) and the average PCT activity (0-0.5) at different concentrations (1e-7 - 1e0 %v/v).

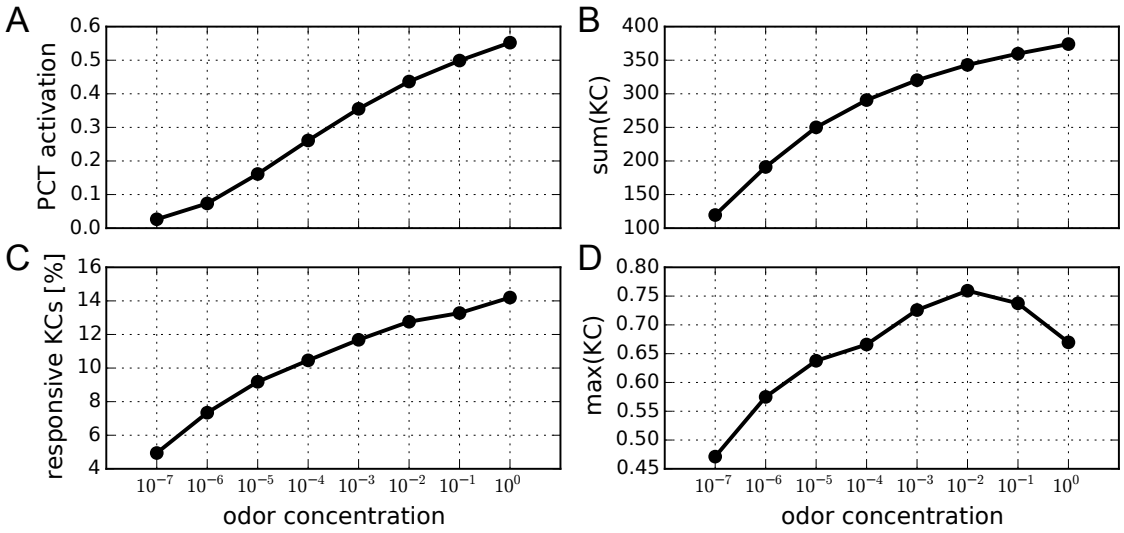

Figure 4.17: **Concentration dependence of KCs and PCTs**. A and B: PCT activation and ratio of active KCs closely follow the predefined heuristic targets. C: Average KC activity monotonically rises with higher concentrations. D: For low concentrations, the highest activity among KCs is roughly proportional to odor concentration. Inhibitory gain control through PCT activity reduces the maximal KC activity for high odor concentrations.

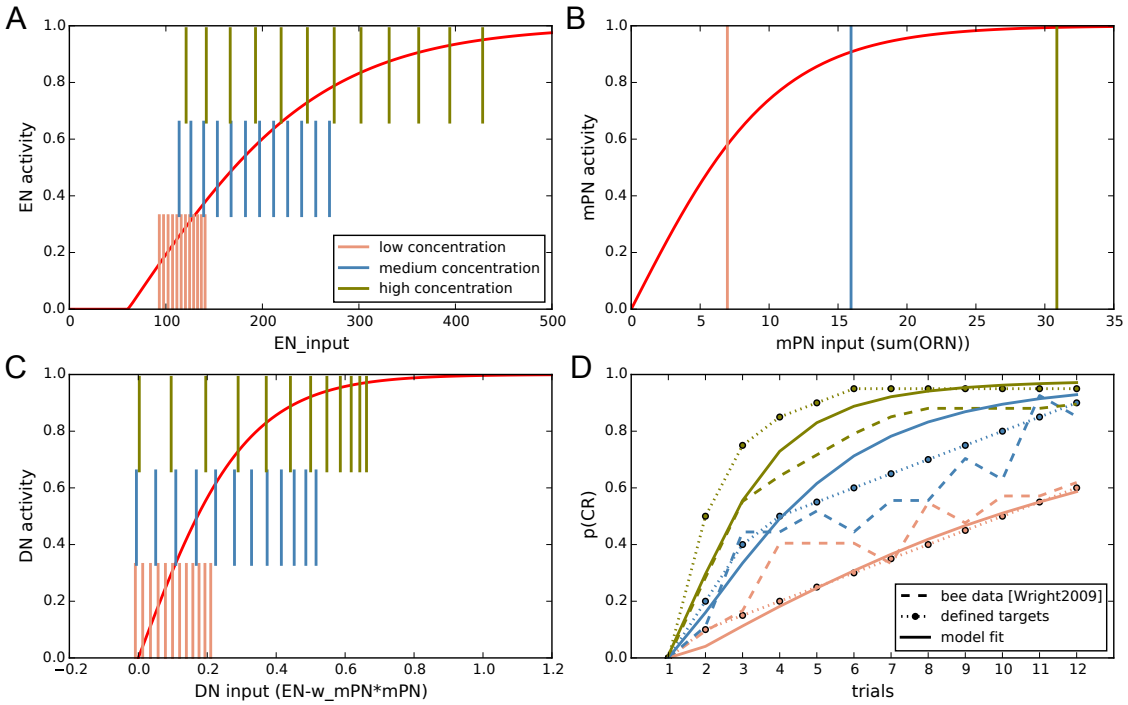

Figure 4.18: **Fitting LH parameters.** Five parameters were fitted in the LH network based on data from (Wright et al., 2009). Prior to this, target were defined as shown in D. A: Input to EN was calculated for absolute conditioning protocols at three different levels of CS concentration (salmon, blue, olive bars). Fitted transfer function of EN (red curve). B: Concentration dependent mPN input (salmon, blue, olive bars). Fitted mPN transfer function (red curve). C: DN input for fitted transfer function of EN and mPN (salmon, blue, olive bars). Fitted DN transfer function (red curve). D: Bee data for absolute conditioning at three levels of concentration (dashed lines). Derived target values for concentration dependent DN activity that equals p(CR) of the model (dots). Best fit based on parameter search for EN transfer function (EN_c, EN_g), LH gain control through mPNs (w^{mPN}, mPN_g), and DN transfer function (DN_g).

Table 4.3: List of all parameters with description, value and further comments.

Parameter	Description	Value [unit]	Comments
	ODOR-REWARD INTEGRATION		
Δt	temporal resolution of reward integration	0.1 [s]	half of the lowest time constant in conditioning protocols (200 ms)
CS onset	start of odor stimulation	0 [s]	defined as t=0s
CS length	duration of odor stimulation	0.2 - 6 [s]	defined in conditioning protocol
US onset	start of sucrose reward	-2 - 14 [s]	defined in conditioning protocol
US length	duration of sucrose reward	1-10 [s]	defined in conditioning protocol, parameter not considered in model
US concentration	molarity of sucrose solution	approx. 1-2 [M]	defined in conditioning protocol, parameter not considered in model
α_E	rate of increase of eligibility trace	0.09 []	fitted to dependence on CS duration [Wright2009c]
τ_E	time constant of decay of eligibility trace	4.3 [s]	fitted to trace conditioning [Szyzska2011]
	ANTENNA		
N_{OR}	number of receptor types	160	defined as the number of receptor types in the honeybee (~163)
OR_m	Hill coefficient of tuning curves	0.26	fitted according to average of coefficients reported in [Sachse2003]
k_0	decimal logarithm of most sensitive receptor	-5	estimated based on most sensitive value reported in [Sachse2003]
k_1	decimal logarithm of least sensitive receptor	12.4	fitted to concentration dependence of compound signal reported in [Sachse2003]
OR_{exp}	exponent determining nonlinear distribution of receptor sensitivities	0.5	fitted to concentration dependence of compound signal reported in [Sachse2003]
s_{XY}	tuning similarity between odor compounds X and Y	0-1	adjusted for some non-absolute conditioning protocols
	ANTENNAL LOBE		
uPN_{sr}	spontaneous rate of uPNs	0.025	heuristic value
uPN_g	broadness of uPN transfer function	9	fitted to transfer function of drosophila PNs [Bhandavat2007]
mPN_g	curvature of mPN transfer function	0.19	fitted to concentration dependent acquisition [Wright2009c] (adjusted values)
$p^{lLN,uPN}$	mean connection probability between lLNs and uPNs	0.25	heuristic value
r_{ex}	ratio of excitatory connections between lLNs and uPNs	0.15	heuristic value
w_{in}^{lLN}	determines initial excitatory weights between lLNs and uPNs	0.03	heuristic value
w_{ex}^{lLN}	determines initial inhibitory weights between lLNs and uPNs	0.075	heuristic value
$w_{in,max}^{lLN}$	maximum of inhibitory weights between lLNs and uPNs	0.2	fitted to differential conditioning data [Pamir2011]
$w_{ex,max}^{lLN}$	maximum of excitatory weights between lLNs and uPNs	0.3	fitted to differential conditioning data [Pamir2011]
$\alpha_{w^{lLN}}$	learning rate for weight changes between lLNs and uPNs	0.25	fitted to differential conditioning data [Pamir2011]
gLN_c	center of gLN transfer function	74	fitted to concentration dependence of uPNs [Sachse2003] (adjusted values)
gLN_g	broadness of gLN transfer function	0.06	fitted to concentration dependence of uPNs [Sachse2003] (adjusted values)
w^{gLN}	initial weights between gLNs and uPNs	0.55	fitted to concentration dependence of uPNs [Sachse2003] (adjusted values)

Table S1 conditued

Parameter	Description	Value [unit]	Comments
	MUSHROOM BODY		
w^{uPN}	weights between uPNs and KCs	0.5	heuristic value
$N_{uPN,KC}$	mean number of uPNs connected to each KC	10	downscaled number of estimated uPNs per KC (50)
N_{KC}	number of KCs	12000	downscaled number of olfactory KCs
KC_c	center of KC transfer function	1.6	fitted to heuristically defined KC concentration dependence
KC_g	broadness of KC transfer function	2.1	fitted to heuristically defined KC concentration dependence
KC_t	threshold of KCs	0.8	fitted to heuristically defined KC concentration dependence
PCT_c	center of PCT transfer function	360	heuristically fitted (10% active KCs with a mean activity of 0.3)
PCT_g	broadness of PCT transfer function	0.015	fitted to heuristically defined KC concentration dependence
w_{init}^{PCT}	initial weight of connections between PCT and w_{uPN}	0.5	arbitrary value
w_{max}^{PCT}	maximal weight of connections between PCT and w_{uPN}	0.6	fitted to latent inhibition [Chandra2010]
w_{min}^{PCT}	minimal weight of connections between PCT and w_{uPN}	0.22	fitted to fast acquisition during absolute conditioning [Pamir2011]
$\alpha^{+}_{w^{PCT}}$	learning rate of weight increase in w_{PCT}	0.75	fitted to extinction dynamics [Stollhoff2005]
$\alpha^{-}_{w^{PCT}}$	learning rate of weight decrease in w_{PCT}	3.25	fitted to fast acquisition during absolute conditioning [Pamir2011]
	LATERAL PROTOCEREBRAL LOBE AND LATERAL HORN		
w^{mPN}	weight between mPN and DN	0.29	fitted to concentration dependence during absolute conditioning [Wright2009c] (adjusted values)
dEN_c	center of dendritic transfer function of EN	0.5	heuristic value
dEN_g	broadness of dendritic transfer function of EN	10	heuristic value
EN_c	center of EN transfer function	61	fitted to concentration dependence during absolute conditioning [Wright2009c] (ajdusted values)
EN_g	broadness of EN transfer function	0.01	fitted to concentration dependence during absolute conditioning [Wright2009c] (adjusted values)
DN_g	broadness of DN transfer function	6.4	fitted to concentration dependence during absolute conditioning [Wright2009c] (adjusted values)

Figures	Reference	ITI [min]	CS	CS conc. [%v/v]	CS dur. [s]	US conc.	US onset [s]	US dur. [s]
4.7A-F	Wright et al. 2009	5	1-hexanol, 1-octanol, 2-octanone	high: ~0.3 medium: ~0.003 low: ~0.00003	0.2, 0.5, 0.8, 1	1.5 M	1	1
4.7G-I	Chandra et al. 2010	5	geraniol, hexanal, 2-octanone, 1-octanal	1.0	4	1.5 M	3	3
4.7J-L	Szyszka et al. 2011	10	1-nonanol, 1-octanol, 1-hexanol, 1-nonanone, 2-heptanone, octanal, citral	0.01	0.5	1.0 M	1-15	3
4.9A, B	Pamir et al. 2001	7	1-octanol, 1-hexanal	1.0	5	30%	3	4
4.9C, D	Yamagata, unpublished Figure 4.9C, D	1.5	1-hexanol, 2-octanol, octanal, 2-octanone	0.01	3	30%	2	3
4.9E-H	Fernandez et al, 2009	6	1-hexanol, 2-octanone	~0.003 (0:10, 9:1)	2	2 M	1	4
4.11A-D	Deisig et al. 2001	8	linalool, 1-hexanol, limonene, 2-octanol	1.0	6	40%	3	3
4.11E, F	Chandra & Smith 1998	n.a.	geraniol, hexanal, 1-hexanol, 2-octanone	1.0	4	1.5 M	3	2

Table 4.4: Details of the classical conditioning protocols that were presented in the results.
Odor concentration were converted to %v/v when given in M.

Chapter 5

General discussion

Chapters 2 - 4 of this thesis are presented in chronological order. While chapter 2 (Helgadottir et al., 2013) was published more than a year ago, chapters 3 and 4 present recent work. This order is also reflected by the content of the chapters. The network model that controls the robot presented in chapter 2 is an early version of the one presented in chapter 4, and was originally used for the investigation of behavioral plasticity observed in honeybees during classical conditioning experiments (Haenicke et al., 2013). The data analysis of chapter 3, in combination with newly available physiological evidence (Sinakevitch et al., 2013), inspired the implementation of the network structure presented in chapter 4. Thus, a new hypothesis was formulated on the basis of the experimental data (chapter 3), which was then implemented as a computational model, and subsequently tested (chapter 4). This, in turn, provided new insights that can be explored by future experiments. Although the robotic implementation could not provide immediate insights into insect olfaction at the stage of research presented in this thesis, such an approach can be advantageous with respect to the complex temporal properties of odor plumes that come with great computational costs in a simulated environment. This aspect, though, prompts an important question: How can the olfactory system of insects, showing intricate temporal dynamics in odor responses across all involved neuropils, be represented in an informative way by an abstract computational model using static patterns?

The computation model presented in chapter 4 was, by necessity, designed on the basis of many assumptions. One of the most drastic ones is the assumption to capture most of the relevant computations in the network while ignoring almost all temporal details concerning the ongoing neural activity. During a single trial, the only temporal relation considered in the model is the relative timing between the conditioned stimulus (CS) and the unconditioned stimulus (US), resolved with a temporal precision of 100 ms. This time step was chosen in relation to the smallest temporal difference between CS-onset and US-onset that was used in the studied conditioning experiments (200 ms). Therefore, the only signal in the network that is computed in a time-resolved way is the synaptic tag Γ_{OA}. Moreover, this tag is only relevant for adjusting the synaptic weights in AL and MB during the ITI. All neural activities in the network during a trial are expressed as a pattern of static rates, which seems like a drastic simplification that

cannot result in the expression of the computational principles at work in the animal. However, there are some arguments that support this level of abstraction in the context of this study.

The central coding element of the model that was to be investigated is the recurrent network structure in the mushroom body (MB). In short, the recurrent inhibitory feedback neuron (PCT) provides both a global signal for gain control as well as bouton-specific modulations for memorizing odor-reward associations, resulting in the expression of a value code in the form of the Kenyon cell (KC) pattern. Here, the assumption is, that default KC responses are low and correspond to a background activity that provides an essential level of excitation in the PCT. Without this minimal inhibitory activity in the MB calyx, KC response patterns towards rewarded odors could not be disinhibited. Rewarded patterns, on the other hand, are assumed to elicit strong responses in a subset of the responsive KCs, at an increased level of PCT activity. This subset of highly active KCs is subsequently detected through the computational properties in the dendritic tree of the MB-extrinsic neuron (EN). Here, the nonlinear properties of individual branches of the EN dendrite result in a transmission of signals from highly active KCs, while suppressing those of "default" KCs. Therefore, the assumption is not that temporal details in the bee brain are not important and can be neglected. On the contrary, small differences in the number of successive spikes within a small time window in a KC are predicted to be a crucial component of the system, and it is assumed that they can be expressed, in this context, in form of a static pattern. In experimental studies, this model prediction has not been confirmed directly, but it fits to the temporal sparseness of KC activity that has been observed in different insect species, including honey bees (Perez-Orive et al., 2002; Szyszka et al., 2005; Ito et al., 2008). For KCs in manduca, it was "found that odor presentations that support associative conditioning elicited only one or two spikes on the odor's onset (and sometimes offset) in each of a small fraction of Kenyon cells" (Ito et al., 2008). According to the model, in such an extreme case of temporal sparseness, it might make a crucial difference whether a KC spikes only ones or twice.

Even if the MB network might be appropriately represented in this way, the question remains whether something is lost computationally, if the neural activity of the antennal lobe (AL) is reduced to a single pattern in each trial. In contrast to KC activity, it has been described that (in locust) "representations are dense, dynamic, and seemingly redundant in the antennal lobe" (Perez-Orive et al., 2002). In addition, diverse dynamics entailing oscillatory synchronization seem to be a common feature across phyla. However, most experimental procedures involve static step stimuli, for which it has been demonstrated repeatedly that PNs show phasic-tonic odor responses with latencies below 200 ms (see chapter 3), LNs were reported to respond even faster (Krofczik et al., 2008). This stereotypic shape of the applied olfactory stimuli is shared among all learning paradigms considered here. Therefore, in the model, odor evoked activities of PNs and LNs are also represented by a static rate pattern, assuming that the initial response in PNs defines the MB input pattern that elicits KC spikes. Exploring the temporal details of

olfactory processing seems only reasonable in the context of stimuli that show detailed temporal dynamics (Geffen et al., 2009).

Apart from this, many physiological details of the bee brain are simplified in the model. One of these is the separation of projections from AL to MB into several antennal lobe tracts (ALT). The two most prominent of these are the lateral and the medial ALT. The corresponding uniglomerular projection neurons (uPN) that connect to the MB via these pathways, derive from two distinct groups of AL glomeruli. These have been suggested to represent parallel olfactory pathways with distinct functional properties (Schmuker et al., 2011; Brill et al., 2013). However, the reported differences between both tracts were of a rather quantitative than qualitative kind. While it is feasible to integrate such a dichotomy into the AL network to provide different degrees of concentration and odorant discrimination in parallel, the number of unknown parameters would increase and would need to be determined heuristically until more knowledge about the concrete connectivity in the AL is available. In addition, based on the randomly generated tuning properties of the given 160 uPNs, two virtual subgroups could be identified and selected to represent these two pathways. Therefore, uPNs in the model are represented in form of a single neuron group.

Related to this, uPNs of both tracts send axons to both the MB and the lateral horn (LH). It has been suggested that the projections to the LH are involved in the control of innate odor-driven behavioral patterns. Under the assumption that, in harnessed honeybees, the extension of the proboscis during conditioning experiments follows internal processes that are largely independent of the innate behavioral repertoire, this detail in connectivity was omitted. Apart from uPNs, the projection pattern of the VUM neurons (cf. section 4.2.1) does also include the LH. Physiological evidence suggests that a few fibers of the medio-lateral ALT, i.e. projections of multiglomerular PNs (mPN), receive octopaminergic input (Sinakevitch et al., 2013). Therefore, it is likely that reward dependent plasticity can be found in the LH. Lacking further knowledge about details in connectivity and function of the LH and its downstream connections, though, the LH network that is proposed in the model includes a gain control mechanism mediated through mPNs that is not subject to learning.

So, which neuropils of the honeybee brain govern associative and non-associative learning? In this thesis, it is suggested that the microglomerular circuitry in the mushroom body calyx plays a key role in a process of distributed plasticity. How different loci of plasticity work together, and how this question might be addressed both experimentally and theoretically, remains to be answered in the future. The computational framework developed in this thesis provides a powerful tool to help in this process.

Bibliography

Abel, R., Rybak, J., and Menzel, R. (2001). Structure and response patterns of olfactory interneurons in the honeybee, Apis mellifera. *The Journal of comparative neurology*, 437(3):363–383.

Ache, B. W. and Young, J. M. (2005). Olfaction: Diverse species, conserved principles. *Neuron*, 48(3):417–430.

Akers, R. and Getz, W. M. (1993). Response of olfactory receptor neurons in honeybees to odorants and their binary mixtures. *Journal of Comparative Physiology A*, 173(2):169–185.

Allison, R. T. (1999). Haematoxylin–from the wood. *Journal of clinical pathology*, 52(7):527–528.

Arena, P., Patané, L., Stornanti, V., Termini, P. S., Zäpf, B., and Strauss, R. (2013). Modeling the insect mushroom bodies: Application to a delayed match-to-sample task. *Neural networks : the official journal of the International Neural Network Society*, 41:202–211.

Bazhenov, M., Huerta, R., and Smith, B. H. (2013). A computational framework for understanding decision making through integration of basic learning rules. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 33(13):5686–97.

Bernardet, U. and Verschure, P. F. M. J. (2010). Iqr: a Tool for the Construction of Multi-Level Simulations of Brain and Behaviour. *Neuroinformatics*, 8(2):113–134.

Bhandawat, V., Olsen, S. R., Gouwens, N. W., Schlief, M. L., and Wilson, R. I. (2007). Sensory processing in the Drosophila antennal lobe increases reliability and separability of ensemble odor representations. *Nature neuroscience*, 10(11):1474–1482.

Bitterman, M. E., Menzel, R., Fietz, A., and Sch?fer, S. (1983). Classical conditioning of proboscis extension in honeybees (Apis mellifera). *Journal of Comparative Psychology*, 97(2):107–119.

Brill, M. F., Rosenbaum, T., Reus, I., Kleineidam, C. J., Nawrot, M. P., and Rössler, W. (2013). Parallel processing via a dual olfactory pathway in the honeybee. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 33(6):2443–56.

Carandini, M. (2012). From circuits to behavior: a bridge too far? *Nature Neuroscience*, 15(4):507–9.

Cassenaer, S. and Laurent, G. (2012). Conditional modulation of spike-timing-dependent plasticity for olfactory learning. *Nature*, 482(7383):47–52.

Chandra, S. B. C. and Smith, B. H. (1998). An analysis of synthetic processing of odor mixtures in the honeybee (Apis mellifera). *The Journal of experimental biology*, 201(Pt 22):3113–21.

Chandra, S. B. C., Wright, G. a., and Smith, B. H. (2010). Latent inhibition in the honey bee, Apis mellifera: Is it a unitary phenomenon? *Animal cognition*, 13(6):805–15.

Chittka, L. and Niven, J. (2009). Are Bigger Brains Better? *Current Biology*, 19(21):R995–R1008.

Darwin, C. (1871). *The descent of man and selection in relation to sex, in Charles Darwin, The origin of species and The descent of man (combined volume)*, volume 5. John Murray, London.

Deisig, N., Giurfa, M., Lachnit, H., and Sandoz, J.-C. (2006). Neural representation of olfactory mixtures in the honeybee antennal lobe. *European Journal of Neuroscience*, 24(April):1161–1174.

Deisig, N., Giurfa, M., and Sandoz, J.-C. (2010). Antennal lobe processing increases separability of odor mixture representations in the honeybee. *Journal of neurophysiology*, 103(4):2185–2194.

Deisig, N., Lachnit, H., and Giurfa, M. (2002). The effect of similarity between elemental stimuli and compounds in olfactory patterning discriminations. *Learning & memory (Cold Spring Harbor, N.Y.)*, 9(3):112–21.

Deisig, N., Lachnit, H., Giurfa, M., and Hellstern, F. (2001). Configural olfactory learning in honeybees: negative and positive patterning discrimination. *Learning & Memory*, pages 70–78.

Deisig, N., Lachnit, H., Sandoz, J.-C., Lober, K., and Giurfa, M. (2003). A modified version of the unique cue theory accounts for olfactory compound processing in honeybees. *Learning & memory (Cold Spring Harbor, N.Y.)*, 10(3):199–208.

Eisthen, H. L. (2002). Why are olfactory systems of different animals so similar? *Brain, Behavior and Evolution*, 59(5-6):273–293.

Erber, J., Masuhr, T., and Menzel, R. (1980). Localization of short-term memory in the brain of the bee, Apis mellifera. *Physiological Entomology*, 5:343–358.

Esslen, J. and Kaissling, K. E. (1976). Zahl und Verteilung antennaler Sensillen bei der Honigbiene (Apis mellifera L.). *Zoomorphologie*, 83(3):227–251.

Felsenberg, J., Gehring, K. B., Antemann, V., and Eisenhardt, D. (2011). Behavioural pharmacology in classical conditioning of the proboscis extension response in honeybees (Apis mellifera). *Journal of visualized experiments : JoVE*, (47):2–5.

Fernandez, P. C., Locatelli, F. F., Person-Rennell, N., Deleo, G., and Smith, B. H. (2009). Associative conditioning tunes transient dynamics of early olfactory processing. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29(33):10191–10202.

Field, D. J. (1994). What Is the Goal of Sensory Coding?

Fonta, C., Sun, X., and Masson, C. (1993). Morphology and spatial distribution of bee antennal lobe interneurones responsive to odours. *Chemical Senses*, 18(2):101–119.

Frost, E. H., Shutler, D., and Hillier, N. K. (2012). The proboscis extension reflex to evaluate learning and memory in honeybees (Apis mellifera): Some caveats. *Naturwissenschaften*, 99(9):677–686.

Ganeshina, O. and Menzel, R. (2001). GABA-immunoreactive neurons in the mushroom bodies of the honey bee: An electron microscopic study. *Journal of Comparative Neurology*, 437(3):335–349.

Geffen, M. N., Broome, B. M., Laurent, G., and Meister, M. (2009). Neural encoding of rapidly fluctuating odors. *Neuron*, 61(4):570–86.

Gerber, B. and Ullrich, J. (1999). No evidence for olfactory blocking in honeybee classical conditioning. *The Journal of experimental biology*, 202 (Pt 13:1839–54.

Giurfa, M. (2003). Cognitive neuroethology: dissecting non-elemental learning in a honeybee brain. *Current Opinion in Neurobiology*, 13(6):726–735.

Giurfa, M. (2007). Behavioral and neural analysis of associative learning in the honeybee: a taste from the magic well. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 193(8):801–824.

Giurfa, M. and Giurfa, M. (2003). The amazing mini-brain: lessons from a honey bee. *Bee World*, 84(1):5–18.

Giurfa, M. and Sandoz, J.-C. (2012). Invertebrate learning and memory: Fifty years of olfactory conditioning of the proboscis extension response in honeybees. *Learning & Memory*, 19(2):54–66.

Groh, C. and Rössler, W. (2011). Comparison of microglomerular structures in the mushroom body calyx of neopteran insects. *Arthropod structure & development*, 40(4):358–67.

Grünewald, B. (1999). Physiological properties and response modulations of mushroom body feedback neurons during olfactory learning in the honeybee, Apis mellifera. *Journal of Comparative Physiology - A Sensory, Neural, and Behavioral Physiology*, 185:565–576.

Hadar, R. and Menzel, R. (2010). Memory formation in reversal learning of the honeybee. *Frontiers in behavioral neuroscience*, 4(December):186.

Haehnel, M. and Menzel, R. (2010). Sensory representation and learning-related plasticity in mushroom body extrinsic feedback neurons of the protocerebral tract. *Frontiers in systems neuroscience*, 4(December):161.

Haenicke, J., Pamir, E., and Nawrot, M. P. (2012). A spiking neuronal network model of fast associative learning in the honeybee.

Haenicke, J., Pamir, E., and Nawrot, M. P. (2013). A computational model of fast associative learning in the honeybee. In *Göttingen Meeting of the German Neuroscience Society, conference talk*.

Hammer, M. (1993). An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. *Nature*.

Hammer, M. (1997). The neural basis of associative reward learning in honeybees. *Trends in neurosciences*, 20(6):245–52.

Hammer, M. and Menzel, R. (1995). Learning and memory in the honeybee. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 15(3 Pt 1):1617–1630.

Hausler, C., Nawrot, M. P., and Schmuker, M. (2011). A spiking neuron classifier network with a deep architecture inspired by the olfactory system of the honeybee. In *2011 5th International IEEE/EMBS Conference on Neural Engineering, NER 2011*, pages 198–202.

Heisenberg, M. (2003). Mushroom body memoir: from maps to models. *Nature reviews. Neuroscience*, 4(4):266–275.

Helgadottir, L. I., Haenicke, J., Landgraf, T., Rojas, R., and Nawrot, M. P. (2013). Conditioned behavior in a robot controlled by a spiking neural network. In *International IEEE/EMBS Conference on Neural Engineering, NER*, pages 891–894.

Hellstern, F., Malaka, R., and Hammer, M. (1998). Backward inhibitory learning in honeybees: a behavioral analysis of reinforcement processing. *Learning & Memory*, 4(5):429–444.

Hildebrand, J. G. and Shepherd, G. M. (1997). Mechanisms of olfactory discrimination: converging evidence for common principles across phyla. *Annual review of neuroscience*, 20:595–631.

Himmelreich, S. and Grünewald, B. (2012). Cellular physiology of olfactory learning in the honeybee brain. *Apidologie*, 43(3):308–321.

H.Mobalegh (2010). HaViMo2 - Image Processing Module.

Hourcade, B., Muenz, T. S., Sandoz, J.-C., Rössler, W., and Devaud, J.-M. (2010). Long-term memory leads to synaptic reorganization in the mushroom bodies: a memory trace in the insect brain? *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(18):6461–5.

Hourcade, B., Perisse, E., Devaud, J.-M., and Sandoz, J.-C. (2009). Long-term memory shapes the primary olfactory center of an insect brain. *Learning & memory (Cold Spring Harbor, N.Y.)*, 16(10):607–15.

Huerta, R. and Nowotny, T. (2009). Fast and robust learning by reinforcement signals: explorations in the insect brain. *Neural computation*, 21(8):2123–2151.

Hussaini, S. A. and Menzel, R. (2013). Mushroom body extrinsic neurons in the honeybee brain encode cues and contexts differently. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 33(17):7154–64.

Indiveri, G., Linares-Barranco, B., Hamilton, T. J., van Schaik, A., Etienne-Cummings, R., Delbruck, T., Liu, S.-C., Dudek, P., Häfliger, P., Renaud, S., Schemmel, J., Cauwenberghs, G., Arthur, J., Hynna, K., Folowosele, F., Saighi, S., Serrano-Gotarredona, T., Wijekoon, J., Wang, Y., and Boahen, K. (2011). Neuromorphic silicon neuron circuits. *Frontiers in neuroscience*, 5(May):73.

Indiveri, G., Stefanini, F., and Chicca, E. (2010). Spike-based learning with a generalized integrate and fire silicon neuron. In *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*, pages 1951–1954.

Ito, I., Ong, R. C.-Y., Raman, B., and Stopfer, M. (2008). Sparse odor representation and olfactory learning. *Nature neuroscience*, 11(10):1177–1184.

Kaulen, P., Erber, J., and Mobbs, P. (1984). Current source-density analysis in the mushroom bodies of the honeybee (Apis mellifera carnica). *Journal of Comparative Physiology A*, 154(4):569–582.

Kenyon, F. (1896). The brain of the bee. A preliminary contribution to the morphology of the nervous system of the arthropoda. *Journal of Comparative Neurology*, 6(3):133–210.

Klinke, I. (2011). *Associative plasticity and context modulation in GABAergic feedback neurons of the mushroom body output in the honeybee (Apis mellifera)*. PhD thesis.

Komischke, B., Sandoz, J.-C., Lachnit, H., and Giurfa, M. (2003). Non-elemental processing in olfactory discrimination tasks needs bilateral input in honeybees. *Behavioural Brain Research*, 145(1-2):135–143.

Krofczik, S., Menzel, R., and Nawrot, M. P. (2008). Rapid odor processing in the honeybee antennal lobe network. *Frontiers in computational neuroscience*, 2(January):9.

Lacher, V. (1964). Elektrophysiologische Untersuchungen an einzelnen Rezeptoren für Geruch, Kohlendioxyd, Luftfeuchtigkeit und Tempratur auf den Antennen der Arbeitsbiene und der Drohne (Apis mellifica L.). *Journal of Comparative Physiology A: Neuroethology,*, 623.

Locatelli, F. F., Fernandez, P. C., Villareal, F., Muezzinoglu, K., Huerta, R., Galizia, C. G., and Smith, B. H. (2012). Nonassociative plasticity alters competitive interactions among mixture components in early olfactory processing. *European Journal of Neuroscience*, (February):n/a–n/a.

Malun, D., Giurfa, M., Galizia, C. G., Plath, N., Brandt, R., Gerber, B., and Eisermann, B. (2002). Hydroxyurea-induced partial mushroom body ablation does not affect acquisition and retention of olfactory differential conditioning in honeybees. *Journal of neurobiology*, 53(3):343–60.

Marter, K., Grauel, M. K., Lewa, C., Morgenstern, L., Buckemüller, C., Heufelder, K., Ganz, M., and Eisenhardt, D. (2014). Duration of the unconditioned stimulus in appetitive conditioning of honeybees differentially impacts learning, long-term memory strength, and the underlying protein synthesis. *Learning & memory (Cold Spring Harbor, N.Y.)*, 21(12):676–85.

Matsumoto, Y., Menzel, R., Sandoz, J.-C., and Giurfa, M. (2012). Revisiting olfactory classical conditioning of the proboscis extension response in honey bees: A step toward standardized procedures. *Journal of Neuroscience Methods*, 211(1):159–167.

Menzel, R. (2012). The honeybee as a model for understanding the basis of cognition. *Nature Reviews Neuroscience*, 13(11):758–768.

Menzel, R. (2014). The insect mushroom body, an experience-dependent recoding device. *Journal of physiology, Paris*.

Menzel, R., Leboulle, G., and Eisenhardt, D. (2006). Small brains, bright minds. *Cell*, 124(2):237–239.

Menzel, R., Manz, G., Menzel, R., and Greggers, U. (2001). Massed and Spaced Learning in Honeybees: The Role of CS, US, the Intertrial Interval, and the Test Interval. *Learning & Memory*, 8(4):198–208.

Meyer, A., Galizia, C. G., and Nawrot, M. P. (2013). Local interneurons and projection neurons in the antennal lobe from a spiking point of view. *Journal of neurophysiology*, 110:2465–74.

Mitchinson, B., Pearson, M. J., Pipe, A. G., and Prescott, T. J. (2011). Biomimetic robots as scientific models: a view from the whisker tip. In Krichmar, J. L. and Wagatsuma, H., editors, *Neuromorphic and Brain-Based Robots*, pages 23–57. Cambridge University Press.

Nawrot, M. M. P., Krofczik, S., Farkhooi, F., and Menzel, R. (2010). Fast dynamics of odor rate coding in the insect antennal lobe. *Arxiv preprint arXiv:1101.0271*, pages 1–17.

Nowotny, T. (2010). Parallel implementation of a spiking neuronal network model of unsupervised olfactory learning on NVidia CUDA. In *Proceedings of the International Joint Conference on Neural Networks*.

Okada, R., Rybak, J., Manz, G., and Menzel, R. (2007). Learning-related plasticity in PE1 and other mushroom body-extrinsic neurons in the honeybee brain. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 27(43):11736–47.

Pamir, E., Chakroborty, N. K., Stollhoff, N., Gehring, K. B., Antemann, V., Morgenstern, L., Felsenberg, J., Eisenhardt, D., Menzel, R., and Nawrot, M. P. (2011). Average group behavior does not represent individual behavior in classical conditioning of the honeybee. *Learning & memory (Cold Spring Harbor, N.Y.)*, 18(11):733–41.

Pamir, E., Szyszka, P., Scheiner, R., and Nawrot, M. P. (2014). Rapid learning dynamics in individual honeybees during classical conditioning. *Frontiers in Behavioral Neuroscience*, 8(September):1–17.

Peele, P., Ditzen, M., Menzel, R., and Galizia, C. G. (2006). Appetitive odor learning does not change olfactory coding in a subpopulation of honeybee antennal lobe neurons. *Journal of comparative physiology. A, Neuroethology, sensory, neural, and behavioral physiology*, 192(10):1083–103.

Pelz, C., Gerber, B., and Menzel, R. (1997). Odorant intensity as a determinant for olfactory conditioning in honeybees: roles in discrimination, overshadowing and memory consolidation. *The Journal of experimental biology*, 200(Pt 4):837–47.

Perez-Orive, J., Mazor, O., Turner, G. C., Cassenaer, S., Wilson, R. I., and Laurent, G. (2002). Oscillations and sparsening of odor representations in the mushroom body. *Science (New York, N.Y.)*, 297(5580):359–365.

Pfeil, T., Grübl, A., Jeltsch, S., Müller, E., Müller, P., Petrovici, M. a., Schmuker, M., Brüderle, D., Schemmel, J., and Meier, K. (2013). Six networks on a universal neuromorphic computing substrate. *Frontiers in Neuroscience*, (7 FEB).

Polsky, A., Mel, B. W., and Schiller, J. (2004). Computational subunits in thin dendrites of pyramidal cells. *Nature neuroscience*, 7(6):621–627.

Quinn, W. G., Harris, W. a., and Benzer, S. (1974). Conditioned behavior in Drosophila melanogaster. *Proceedings of the National Academy of Sciences of the United States of America*, 71(3):708–712.

Rath, L., Giovanni Galizia, C., Szyszka, P., and Galizia, C. G. (2011). Multiple memory traces after associative learning in the honey bee antennal lobe. *The European journal of neuroscience*, 34(2):352–60.

RobotShop Inc. (2012). DfRobotShop Rover User Guide, v2.1.

Rost, T., Ramachandran, H., Nawrot, M. P., and Chicca, E. (2013). A neuromorphic approach to auditory pattern recognition in cricket phonotaxis. *Ieee*, 1:1–4.

Rybak, J. and Menzel, R. (1993). Anatomy of the mushroom bodies in the honey bee brain: The neuronal connections of the alpha-lobe. *Journal of Comparative Neurology*.

Rybak, J. and Menzel, R. (1998). Integrative Properties of the Pe1 Neuron, a Unique Mushroom Body Output Neuron. *Learning & Memory*, 5(1):133–145.

Rybak, J. and Menzel, R. (2010). Mushroom body of the honeybee. In *Handbook of Brain Microcircuits*.

Sachse, S. and Galizia, C. G. (2003). The coding of odour-intensity in the honeybee antennal lobe: local computation optimizes odour representation. *European Journal of Neuroscience*, 18(8):2119–2132.

Sachse, S., Galizia, C. G., Brandstaetter, A. S., Kleineidam, C. J., and Riffell, J. A. (2002). Role of inhibition for temporal and spatial odor representation in olfactory output neurons: a calcium imaging study. *Journal of neurophysiology*, 87:1106–1117.

Sandoz, J.-C. (2011). Behavioural and neurophysiological study of olfactory perception and learning in honeybees. *Frontiers in Systems Neuroscience*, 5(December):98.

Scheiner, R., Page, R. E., and Erber, J. (2001). The effects of genotype, foraging role, and sucrose responsiveness on the tactile learning performance of honey bees (Apis mellifera L.). *Neurobiology of learning and memory*, 76(2):138–150.

Scheiner, R., Page, R. E., and Erber, J. (2004). Sucrose responsiveness and behavioral plasticity in honey bees (Apis mellifera). *Apidologie*, 35(2):133–142.

Scheunemann, L., Jost, E., Richlitzki, A., Day, J. P., Sebastian, S., Thum, A. S., Efetova, M., Davies, S.-a., Schwärzel, M., and Schwa, M. (2012). Consolidated and labile odor memory are separately encoded within the Drosophila brain. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 32(48):17163–71.

Schmuker, M., Pfeil, T., and Nawrot, M. P. (2014). A neuromorphic network for generic multivariate data classification. *Proceedings of the National Academy of Sciences of the United States of America*, 111(6):2081–6.

Schmuker, M., Yamagata, N., Nawrot, M. P., and Menzel, R. (2011). Parallel representation of stimulus identity and intensity in a dual pathway model inspired by the olfactory system of the honeybee. *Frontiers in neuroengineering*, 4(December):17.

Schröter, U., Malun, D., and Menzel, R. (2007). Innervation pattern of suboesophageal ventral unpaired median neurones in the honeybee brain. *Cell and tissue research*, 327(3):647–67.

Sinakevitch, I. T., Smith, A. N., Locatelli, F., Huerta, R., Bazhenov, M., and Smith, B. H. (2013). Apis mellifera octopamine receptor 1 (AmOA1) expression in antennal lobe networks of the honey bee (Apis mellifera) and fruit fly (Drosophila melanogaster). *Frontiers in systems neuroscience*, 7(October):70.

SmartProjects - Italy (2005). Arduino Uno Datasheet.

Smith, B. H. (1998). Analysis of interaction in binary odorant mixtures. *Physiology & behavior*, 65(3):397–407.

Smith, B. H. and Cobey, S. (1994). The olfactory memory of the honeybee Apis mellifera. II. Blocking between odorants in binary mixtures. *The Journal of experimental biology*, 195:91–108.

Soltic, S. and Kasabov, N. (2011). *System and Circuit Design for Biologically-Inspired Intelligent Learning*. IGI Global.

Stollhoff, N., Menzel, R., and Eisenhardt, D. (2005). Spontaneous recovery from extinction depends on the reconsolidation of the acquisition memory in an appetitive learning paradigm in the honeybee (Apis mellifera). *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 25(18):4485–92.

Strube-Bloss, M. F., Nawrot, M. P., and Menzel, R. (2011). Mushroom Body Output Neurons Encode Odor-Reward Associations. *Journal of Neuroscience*, 31(8):3129–3140.

Szyszka, P., Demmler, C., Oemisch, M., Sommer, L., Biergans, S., Birnbach, B., Silbering, a. F., and Galizia, C. G. (2011). Mind the Gap: Olfactory Trace Conditioning in Honeybees. *Journal of Neuroscience*, 31(20):7229–7239.

Szyszka, P., Ditzen, M., Galkin, A., Galizia, C. G., and Menzel, R. (2005). Sparsening and temporal sharpening of olfactory representations in the honeybee mushroom bodies. *Journal of neurophysiology*, 94(5):3303–3313.

Szyszka, P., Galkin, A., and Menzel, R. (2008). Associative and non-associative plasticity in kenyon cells of the honeybee mushroom body. *Frontiers in systems neuroscience*, 2(June 2008):3.

Takeda, K. (1961). Classical conditioned response in the honey bee.

Weiner, J. (1999). *Time, Love, Memory: A Great Biologist and His Quest for the Origins of Behavior*. Vintage Books.

Wessnitzer, J., Young, J. M. J. J. M., Armstrong, J. D., and Webb, B. (2012). A model of non-elemental olfactory learning in Drosophila. *Journal of computational neuroscience*, 32(2):197–212.

Witthöft, W. (1967). Absolute anzahl und verteilung der zellen im him der honigbiene. *Zeitschrift für Morphologie der Tiere*, 61(1):160–184.

WIZnet Co. Ltd. (2011). WizFi210/220 User Manual2, v1.0.

Wright, G. a. (2004). Different Thresholds for Detection and Discrimination of Odors in the Honey bee (Apis mellifera). *Chemical Senses*, 29(2):127–135.

Wright, G. a., Carlton, M., and Smith, B. H. (2009). A honeybee's ability to learn, recognize, and discriminate odors depends upon odor sampling time and concentration. *Behavioral Neuroscience*, 123(1):36–43.

Wright, G. a., Thomson, M. G. a., and Smith, B. H. (2005). Odour concentration affects odour identity in honeybees. *Proceedings. Biological sciences / The Royal Society*, 272(1579):2417–22.

Yamagata, N., Schmuker, M., Szyszka, P., Mizunami, M., and Menzel, R. (2009). Differential odor processing in two olfactory pathways in the honeybee. *Frontiers in systems neuroscience*, 3(December):16.

Young, J. M., Wessnitzer, J., Armstrong, J. D., and Webb, B. (2011). Elemental and non-elemental olfactory learning in Drosophila. *Neurobiology of learning and memory*, 96(2):339–352.