Spike-Timing Dependent Plasticity and Synchronous Oscillations in an Invertebrate Olfactory System

Thesis by
Stijn Cassenaer

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

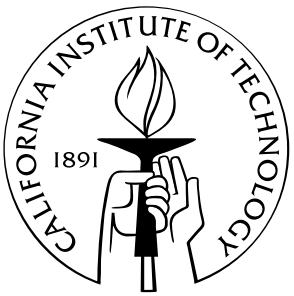

California Institute of Technology
Pasadena, California

2008
(Defended September 6, 2007)

Aan mijn ouders

Acknowledgements

I wish to thank my advisor, Gilles Laurent, for creating a fascinating research program that has attracted great people with whom I've had the benefit to interact; for offering me difficult and very rewarding research projects; for valuable insights; and for appreciation of electrophysiological beauty. I also wish to thank the other members of my Thesis Committee: Mark Konishi, Pietro Perona and Thanos Siapas, for valuable support over many years. Scientifically, I am very much indebted to several generations of the entire Laurent lab. I am particularly grateful to Vivek Jayaraman, for many reasons, but perhaps most for a very inspiring love of science and demonstration of excellence, and for strongly encouraging independent thinking. I am also particularly grateful to Maria Papadopoulou, for the most stimulating and satisfying scientific discussions, and for a perfect mixture of criticism, encouragement and helpfulness.

On a personal level, I am very grateful for the friendship of many generations of the Laurent lab. There are several people to whom I am particularly grateful for intense friendship during my time at Caltech: Scott Detmer, Dan Gold, Vivek Jayaraman, Ofer Mazor, Anusha Narayan, Maria Papadopoulou, Melanie Pribisko-Yen, and Glenn Turner. I am grateful to Vivek Jayaraman for helping me navigate considerable personal difficulties, and for impactful philosophical insights. I am also very grateful to Ofer Mazor for teaching me how to properly ride a bicycle,

and for helping me get started in the Laurent lab. To Javier Perez-Orive for many movie discussions; to Cindy Chiu, Ben Rubin and Laurent Moreaux for exemplifying beauty in method; to Roni Jortner for extraordinary stories; to Laurent Moreaux for much entertainment and culinary common-ground; to Mark Stopfer for demonstrating experimental beauty and for understanding serpentine matters; to Rachel Wilson for scientific vigor; to Mattias Westman for snow on the roofs in Moskou; to Alex Backer for the scent of darkness; to Sarah Farivar, Ofer Mazor, Javier Perez-Orive, Ben Rubin, and Kai Shen for Guacamole fellowship; to Glenn Turner for the Skating; to Anusha Narayan for several lifetimes' supply of puns. And I want to thank Maria Papadopoulou for very many unspecified reasons. I owe very much to my parents and to my family for love, encouragement and opportunities.

Abstract

Sensory systems neuroscience aims to study how patterns of neural activity represent stimuli of the outside world. To this end, the present work addresses how olfactory stimuli are represented by three successive layers in the locust olfactory system. Activation by an odorant of primary sensory neurons in the antenna gives rise to broadly distributed, oscillatory spatiotemporal activity patterns across the antennal lobe (AL). This is in marked contrast to the representation in the mushroom body (MB), where Kenyon cells (KCs) respond very sparsely and very briefly. In the AL, an odor gives rise to a particular trajectory through Projection Neuron (PN) phase space, with individual timepoints representing different aspects of the stimulus; in the MB, very small subsets of KCs respond selectively at particular timepoints along this trajectory. Two mechanisms are identified that contribute to the sparsening across the two structures: an intrinsic voltage dependence in the KCs, which gives rise to a superlinear response to synchronous inputs, and a canonical network motif, feedforward inhibition, which diminishes the KC response to nonsynchronous excitatory inputs. From a decoding perspective, this makes the oscillation cycle the relevant timestep of the AL trajectories, and it demonstrates a role for synchronous oscillations in sensory networks. While broad activation of the AL promotes extensive local interactions, giving rise to dynamic representations and enabling multiple features to be extracted, the sparse representation after

decoding by KCs likely facilitates the storage of relevant patterns in memory.

A subset of MB extrinsic neurons with dendrites densely invading the β-lobe (βLNs) is well placed to decode the KCs' sparse responses. The synapses formed by KCs onto these cells are powerful and undergo Hebbian spike-timing dependent plasticity (STDP) on a timescale similar to the synchronous oscillations generated in the AL (and propagated through the MB). STDP has a homeostatic effect on the firing phase of βLNs by fine-tuning the strength of KC-βLN synapses, contributing to tight locking among subsets of βLNs during odor stimulation and facilitating the flow of synchronous information.

The facilitation of tight synchrony among βLNs by STDP further ensures that different odor features computed and formatted as a function of cycle number by the AL, and represented by the sparse representations of KCs, remain segregated between LFP oscillation cycles. This segregation is also sustained by phase-locked feedforward inhibition onto βLNs, which restricts the window of integration for inputs from KCs, and is found to be due to neighboring βLNs of the same class. The implications of the resultant competition among βLNs due to this inhibition, and particularly its interaction with STDP at the KC-βLN synapse are addressed with a network model. The results are considered within the context of the circuit in which the KC-βLN network is embedded, and a cycle-specific mechanism for learning an arbitrary subset of the odor features computed in the AL is proposed.

Contents

List of Figures

Chapter 1

Introduction

1.1 Early Olfactory Coding

Olfactory information enters the nervous system when odorant molecules bind to olfactory receptors (ORs, class A G-protein coupled receptors) embedded in the dendritic membrane of olfactory receptor neurons (ORNs, Buck and Axel, 1991; Buck, 1996). ORNs are located in the antennae of insects and in the olfactory epithelia of vertebrates. Individual ORNs express one or very few ORs (probably one in the case of mammals (Ressler et al., 1993; Vassar et al., 1993; Laurent et al., 2001; Mombaerts, 2004b); for *Drosophila*, most ORNs express one OR, but a few ORNs express two or three (Vosshall et al., 1999; Hallem et al., 2004; Goldman et al., 2005; Couto et al., 2005)). In the antenna, as in the olfactory epithelium, all ORNs that express the same OR are localized to a particular area, representing one of a small number of large overlapping zones. Within a given zone, ORNs of different types are distributed seemingly randomly (Ressler et al., 1993; Weth et al., 1996; de Bruyne et al., 2001).

In insects, ORNs send axons, via the antennal nerve (AN), to the antennal lobe (AL), where they contact the principal neurons of the AL, projection neurons (PNs,

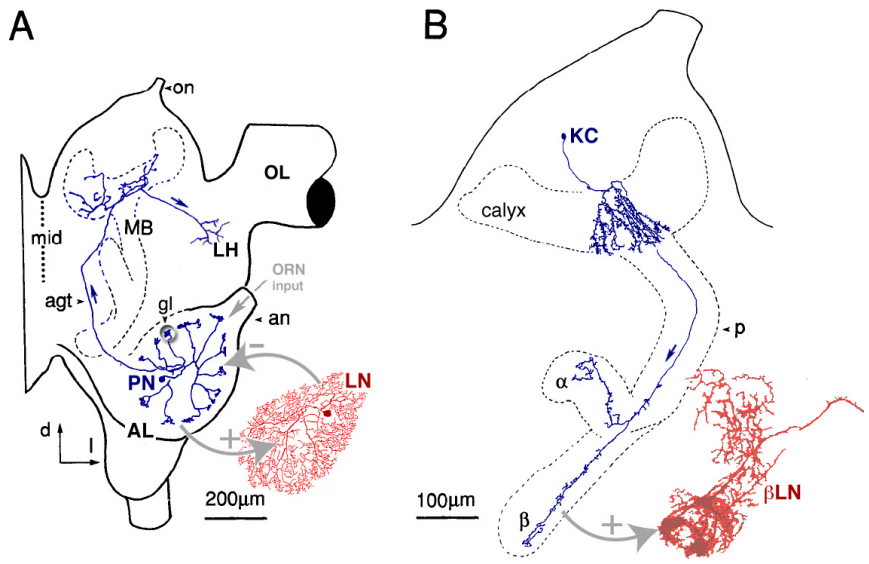

Figure 1.1. Locust Olfactory Anatomy. AL, antennal lobe; LH, lateral horn; MB, mushroom body; OL, optic lobe; agt, antennal-glomerular tract; an, antennal nerve; gl, glomerulus; on, ocellar nerve; p, pedunculus; βLN, β-lobe neuron; KC, Kenyon cell; LN, local neuron; ORN, olfactory receptor neuron; PN, projection neuron; d, dorsal; l, lateral; mid, midline. Adapted from Laurent and Naraghi (1994); MacLeod and Laurent (1996).

cholinergic), as well as local neurons (LNs, GABAergic, but, in *Drosophila*, some also cholinergic (Shang et al., 2007)). The neuropil where these synapses are formed is organized into glomeruli (Figure 1.1A). PNs send axons out of the AL, via the antennal glomerular tract (AGT), and synapse onto the dendrites of Kenyon cells (KCs) in the calyx of the mushroom body (MB). Beyond the MB, the bifurcating axons of PNs also target the Lateral Horn (LH), where they contact inhibitory neurons (LHIs) (Laurent et al., 2001).

This architecture resembles a simplified version of the mammalian olfactory bulb (OB): ORNs send axons, via the olfactory nerve (ON), to the OB, where they contact the principal neurons of the OB, mitral and tufted (M/T, glutamatergic) cells, as well as periglomerular (PG, GABAergic and/or dopaminergic) cells. The neuropil where these synapses are formed is organized into glomeruli. M/T cells

send axons out of the OB, via the lateral olfactory tract (LOT), and synapse onto the apical dendrites of pyramidal cells in Layer Ia of the olfactory cortex (anterior olfactory cortex and piriform cortex). Beyond the olfactory cortex, the axons of the M/T cells also target the prefrontal cortex, the olfactory turbercle, the amygdala and the entorhinal cortex. In addition to sending an apical dendrite into the glomerular layer, M/T cells have several secondary dendrites in the external plexiform layer (EPL), where they contact granule cells (GCs, GABAergic); they also extend axon collaterals within the granule cell layer (GCL) (Shepherd, 2004; Wilson and Mainen, 2006). Table 1.1 lists olfactory neuron numbers for a few prominent species.

Organism	ORs	ORNs	Glomeruli	LNs (or P/GCs)	PNs (or MCs)	KCs (or PCs)
Fruitfly	62	1,300	50	100	150-200	2,500
Honeybee	150	60,000	156-166	4,000	800	170,000
Locust		90,000	1,000	300	830	50,000
Mouse	1200	2-20x10^6	1,800	2.5-4.5x10^6	38,000	>380,000
Zebrafish	112	10^6	80	2.5-7.8x10^4	350 to 650	

Table 1.1. Olfactory Neuron Numbers (per hemisphere)

Despite the intermingling of different ORN types within (particular zones of) the antenna and olfactory epithelium, once their axons enter the AL or OB, all ORNs of a particular type target the same glomerulus (typically, a bilateral pair in *Drosophila* and a unilateral pair in mice) with great precision (Stocker et al., 1990; Vassar et al., 1994; Mombaerts et al., 1996; Bozza and Kauer, 1997; Gao et al., 2000; Couto et al., 2005). Such an arrangement would suggest a rather straightforward, labeled-line coding scheme. Given that ORNs generally express only one OR, when an odorant accesses the antenna or epithelium, it activates the ORNs that express its cogent receptor, which, in turn, activates its specific glomerulus in the AL or OB. Given that mouse M/T cells and *Drosophila* PNs send a primary neurite to only

one glomerulus, the activity of a given PN or M/T cell type (defined by glomerular projection) would signify the presence of that particular odor. This would not be taking into account, however, the fact that individual odorants can activate a broad set of OR(N)s, and that individual OR(N)s can be activated by multiple, even chemically dissimilar odorants (Bozza and Kauer, 1997; Malnic et al., 1999). In addition, even in species with monoglomerular M/T cells or PNs (as of course in multiglomerular species such as locust and zebrafish), there are interactions between glomeruli (Meister and Bonhoeffer, 2001; Shepherd, 2004). Such interactions could, in principle, have the effect of either broadening or sparsening the principal neurons' odor tuning, and both have been reported (Friedrich and Laurent, 2004; Wilson et al., 2004; Shang et al., 2007; Olsen et al., 2007) and (Yokoi et al., 1995; Friedrich and Laurent, 2004). The odor tuning of ORNs, as well as of M/T cells and PNs, has been examined in electrophysiology and imaging experiments, and, in all cases, some cells are quite promiscuous and others very selective[1] (Malnic et al., 1999; de Bruyne et al., 2001; Hallem et al., 2004). For PNs in the insect AL, as well as M/T cells in lower vertebrates, the evidence is fairly unequivocal that tuning is generally broad (Joerges et al., 1997; Friedrich and Korsching, 1997; Perez-Orive et al., 2002; Carlsson et al., 2002; Wilson et al., 2004; Mazor and Laurent, 2005; Shang et al., 2007; Olsen et al., 2007) except for the report by Wang et al. (2003). For M/T cells in mammals, there are many studies suggesting generally broad tuning (Adrian, 1942; Adrian, 1950; Leveteau and MacLeod, 1966; Stewart et al., 1979; Cinelli and Kauer, 1992; Motokizawa, 1996; Rubin and Katz,

[1]This is the case for the general olfactory system. Both insects and mammals also employ an olfactory subsystem comprised of highly selectively tuned cells for detecting a specific odorant (e.g., a pheromone), which typically gives rise to an innate behavioral response. This subsystem can consist of a single class of ORNs (Suh et al., 2004), and/or an individual glomerulus (Christensen et al., 1995; Hansson et al., 2003; Suh et al., 2004), or separate structures altogether, i.e., the vomeronasal organ and accessory olfactory bulb (Mombaerts, 2004a).

1999; Lehmkuhle et al., 2006), but also two recent studies (Rinberg et al., 2006; Davison and Katz, 2007) reporting predominantly sparse representations. Rinberg et al. (2006) furthermore describe sparsening of M/T cells' responses following recovery from anaesthesia. It is possible that these studies uncover a fundamental difference in terms of density of representation between the OB and AL and that very different coding schemes and information processing strategies are implemented by the two structures. It seems worth noting, however, that, even if the fraction of the population responding to an odor were on average an order of magnitude smaller (e.g. 5% of M/T cells in the mouse OB compared to 50% of PNs in the locust AL), the actual number responding would still amount to approximately five times more M/T cells in the OB (than PNs in the AL) participating in the odor's representation (Table 1.1). Also, Rinberg et al. (2006) report a dramatic increase in baseline firing rate that accompanies the apparent sparsening in the awake state, which, as the authors point out, could allow for a temporal coding dimension not evaluated in their study (nor in Davison and Katz (2007)). Doing so might reveal broader M/T cell participation in odor coding than would be estimated from firing rates alone.

The temporal aspect of the olfactory code has received considerable attention in invertebrates and vertebrates alike, and temporal patterning of odor responses are observed in both uni- and multiglomerular principal neurons of the AL and OB (Macrides and Chorover, 1972; Burrows et al., 1982; Meredith, 1986; Wellis et al., 1989; Buonviso et al., 1992; Yokoi et al., 1995; Motokizawa, 1996; Laurent et al., 1996; Spors and Grinvald, 2002; Perez-Orive et al., 2002; Wilson et al., 2004; Mazor and Laurent, 2005). This occurs, whether odor pulses are square or intermittent (Brown et al., 2005; Broome et al., 2006), in contrast to the relative simplicity of

ORN responses (Duchamp-Viret et al., 2000; Duchamp-Viret et al., 1999; de Bruyne et al., 2001; Friedrich and Laurent, 2001; Friedrich and Laurent, 2004; Wilson et al., 2004; Wilson and Laurent, 2005) except (Spors et al., 2006), and is thought to arise predominantly from the interactions within the AL and OB (Laurent et al., 2001; Laurent, 2002). It can be characterized as having relatively fast (10s of ms) and slow (100s of ms) components. The fast component is oscillatory ($\beta - \gamma$ range) and essentially due to local inhibitory interactions, mediated by $GABA_A$ receptors (MacLeod and Laurent, 1996; MacLeod et al., 1998; Wehr and Laurent, 1999; Stopfer and Laurent, 1999; Perez-Orive et al., 2002; Friedrich et al., 2004). In the OB electrical synapses play a role as well (Friedman and Strowbridge, 2003; Lowe, 2003). The slow component derives from local excitatory and inhibitory (including $GABA_B$ receptor-mediated) interactions (Hamilton and Kauer, 1989; Wellis et al., 1989; Buonviso et al., 1992; Isaacson and Strowbridge, 1998; Aroniadou-Anderjaska et al., 1999; Friedrich and Laurent, 2001; Luo and Katz, 2001; Margrie et al., 2001; Urban and Sakmann, 2002; Cang and Isaacson, 2003; Hayar et al., 2004a; Wilson and Laurent, 2005; Shang et al., 2007; Olsen et al., 2007), including presynaptic inhibition of ORN terminals (Wachowiak and Cohen, 1999; Aroniadou-Anderjaska et al., 2000; Isaacson and Vitten, 2003; Wachowiak et al., 2005), some of which is dopamine-mediated (Wachowiak and Cohen, 1999; Ennis et al., 2001).

In addition to the temporal features that are generated intrinsically, there are also fluctuations in odorant exposure due to antennal flicking (insects), coughing (fish) and respiration (terrestrial vertebrates) (Schaefer and Margrie, 2007), which modulate neural activity accordingly. This has been particularly explored in rodents, where breathing imposes a θ rhythm on M/T cells' activity. It appears that the θ rhythm is to some extent intrinsic to the M/T cells, and that the population is

entrained by respiration (Hayar et al., 2004b). M/T cells' spiking latency relative to this sniff cycle has been proposed as a code for odor identity: an odor evokes a specific pattern of latencies across multiple M/T cells (Margrie and Schaefer, 2003; Schaefer and Margrie, 2007), as described previously by Hopfield (1995). Furthermore, with increasing concentration, M/T cells' firing latencies decrease, but the relative order of different cells' firing remains the same, thus endowing concentration invariance, as in the Hopfield (1995) model.

Such a coding scheme is quite distinct from the way odors are encoded in the AL. When an odor is presented, PNs fire at a particular average phase relative to the oscillatory population response (measured as a local field potential, LFP), and there is no apparent information, about odor identity (Laurent et al., 1996) or concentration (Stopfer et al., 2003), contained in the phase of PN spikes. Instead, PNs fire at distinct subsets of cycles of the LFP in a PN- and odor-specific manner (Wehr and Laurent, 1996), and, when considered as an ensemble, the activity pattern varies continuously across odor concentration (Stopfer et al., 2003). The significance of PN spikes locking to the LFP (indeed their synchrony, which gives rise to it), their cycle-specificity, and their broad tuning becomes apparent when considered from the perspective of their downstream decoders, the KCs. This is the subject of Chapter 2. Briefly, KCs respond to odors quite sparsely, despite their inputs' broad tuning, because of intrinsic and network properties which make them behave as coincidence detectors (Perez-Orive et al., 2002; Perez-Orive et al., 2004). This tunes the KCs selectively to synchronous inputs, a result confirmed by the observation that the KC population firing profile matches well the profile of PN phase-locking strength (Mazor and Laurent, 2005). As mentioned, PNs fire at a subset of cycles in a PN- and odor-specific manner, spanning

just a few cycles on average (Mazor and Laurent, 2005), reflecting the rate at which the activity of the PN population vector is updated (Mazor and Laurent, 2005). KCs appear to be connected to approximately half of the PNs, which, individually, have a very small impact (80 microvolts on average) relative to the KC firing threshold (Jortner et al., 2007). The model, then, is that KCs sample a large, evolving activity vector, and fire a spike only when the right combination presents itself (Broome et al., 2006). This does not happen very often (because many of the connected PNs must fire simultaneously) but nevertheless the number of combinations that could make an individual KC fire is quite large (Jortner et al., 2007). Different KCs see their preferred PN ensembles activate at different times after the odor is encountered and experiments in zebrafish have provided an example of what such different time-points could functionally correspond to (Friedrich and Laurent, 2001). These experiments examined how odor representations in the OB evolve over a time-course of 100s of ms to seconds relative to odor onset. The main finding is that representations become de-correlated over time, such that initially similar activity patterns evoked by (chemically related) odorants, become dissimilar. From a functional point of view, early in the response, activity in the OB represents information about the class to which an odorant belongs, whereas the representation at a later point reflects the identity of the particular odorant. Consequently, cells decoding OB activity would perform either stimulus classification or identification, depending on whether their preferred ensemble occurs early or late.

Given that olfactory stimuli evolve over a relatively slow time-scale, compared to their visual or auditory counterparts, time itself has been proposed as a coding variable. However, the above model of piecewise decoding, together with the de-correlation results from the zebrafish experiments, suggests that the slow spa-

tiotemporal patterns described above reflect a transformation carried out by the encoding structure, and are not explicitly decoded as a sequence. As such, time is not used[2] as a coding variable per se; instead, the olfactory system takes advantage of the slowness of its stimuli by computing very close to the periphery. It uses its encoding units to compute and subsequently represent different aspects of the same stimulus at different times.

1.2 Associative Memory

1.2.1 Synaptic Plasticity

Learning can be defined as a change in behavior (or knowledge) based on previous experience. In animals it is generally assumed to require a physical change in the brain, and the predominant theory is that memory is stored in the strength of synapses between neurons. An early postulate of *how* synapses are modified was put forth by Donald Hebb:

> *"When an axon of cell A is near enough to excite cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased* (Hebb, 1949)."

Experimental evidence for a long lasting change in synaptic efficacy resulting from such co-activation was first described[3] by Bliss and Lømo in 1973. This effect, known as long term potentiation (LTP), has been studied by recording changes in

[2] It is, of course, conceivable that such sequence decoding occurs elsewhere in the system, e.g., of PN activity by cells in the LH, or of KCs by MB extrinsic cells.

[3] Although reports of repetitive activation resulting in synaptic efficacy enhancement on a shorter timescale were published almost 20 years earlier (Cragg and Hamlyn, 1955; Andersen, 2003)

excitatory postsynaptic potentials (EPSPs, or - currents, EPSCs) extracellularly or intracellularly, in response to high frequency stimulation of afferent fibers; or alternatively, by low frequency stimulation paired with postsynaptic depolarization (by current injection). In fact, a multitude of induction protocols have been found to produce LTP (Bliss and Collingridge, 1993). A particular subset of stimulation paradigms has also been identified which can persistently alter synaptic efficacy in the opposite direction, i.e., long term depression (LTD).

A molecule much praised for virtually uniquely endowing associativity by detecting the paired activity of pre- and postsynaptic cells, was identified in the 1980s as the *N*-methyl d-aspartate receptor (NMDAR, reviewed by Collingridge and Bliss, 1987). The NMDA receptor has a dual dependence for channel-opening, on both presynaptic release of glutamate and postsynaptic membrane depolarization. When the postsynaptic membrane is hyperpolarized, Mg^{2+} blocks the channel's pore, and only when this Mg^{2+} block is relieved by depolarization, *and* glutamate is bound to the receptor, will the channel become significantly permeable to cations. The resultant postsynaptic increase, particularly of calcium, gives rise to a sustained modification of synaptic efficacy (as discussed further below).

During the first two decades following the discovery of LTP and LTD, the requirement for induction was generally considered to be correlation of pre- and postsynaptic cell activity, and the precise timescale of co-activation was typically not addressed. Furthermore, since potentiation and depression were induced by protocols that largely did not evaluate whether postsynaptic cells reached threshold, and because it was unclear how it should be relevant for synapses out in the dendrites, the role of the postynaptic action potential (AP) in synaptic plasticity was not appreciated. The observation by Stuart and Sakmann (1994) that APs

can propagate back into the dendrites of pyramidal neurons, however, led to the recognition of AP timing as an associative signal for Hebbian plasticity (Magee and Johnston, 1997; Markram et al., 1997). It was realized that the back-propagating action potential could be thought of as a global (i.e., cell-wide) signal of postsynaptic activity, broadcast to all the cell's synapses, including, most relevantly, those that have just received, or are about to receive, input from presynaptic cells.

Spike-Timing Dependent Plasticity (STDP) is generally credited to Markram et al.(1997), as well as Bi and Poo (1998), and refers to persistent changes in synaptic strength that result from paired pre- and postsynaptic activity on a millisecond timescale. Specifically, in the most commonly described Hebbian form, a presynaptic AP *following* a postsynaptic AP within a particular time window results in a *decrease* of synaptic efficacy. Reversing the order of pre- and postsynaptic APs within this window *enhances* the strength of the synapse (Figure 1.2). Outside the narrow (50-100ms) window, synaptic weights are unaffected; within it, the polarity of the change can switch abruptly. STDP has been described in many vertebrate

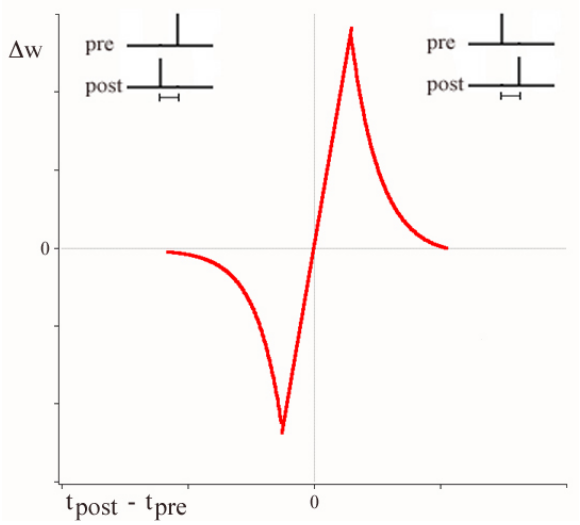

Figure 1.2. Hebbian STDP: synaptic weight changes (Δw) as a function of the relative timing (Δt) of pre- and postsynaptic APs.

species and preparations (Dan and Poo, 2006), including rat (hippocampal culture (Bi and Poo, 1998; Li et al., 2004), hippocampal slice (Debanne et al., 1998; Nishiyama et al., 2000; Lin et al., 2003), neocortical slice (Markram et al., 1997), barrel cortical slice (Egger et al., 1999; Feldman, 2000), visual cortical slice (Sjostrom et al., 2001; Froemke et al., 2005)), mouse (brain stem slice (Tzounopoulos et al., 2004)), zebra finch (brain slice (Boettiger and Doupe, 2001)), electric fish (electric sensory lobe slice (Bell et al., 1997)), and *Xenopus* (retino-tectal system *in vivo* (Zhang et al., 1998)). Different dependencies, in terms of synaptic weight changes as a function of the relative timing of pre- and postsynaptic APs, have been described for different cell types. Figure 1.3 (adapted from Abbott and Nelson (2000)) summarizes data for five different STDP curves.

In addition to direct measurements of synaptic strength (by presynaptic stimulation and intracellular recording of postsynaptic potentials), several studies have assessed, indirectly, the effect of STDP at the circuit level (Dan and Poo, 2006). Generally, this is done, *in vivo*, by sensory stimulation and evaluation of changes in representation consistent with STDP at interposed synapses. Such manipulations of sensory representation have been carried out in cat visual cortex, by ms-timescale pairing of visual and electrical stimulation, and measuring changes in orientation tuning with intrinsic imaging (Schuett et al., 2001); and by pairing gratings at different orientations, and measuring tuning changes with extracellular single unit recordings (Yao and Dan, 2001; Felsen et al., 2002). Notably, plasticity experiments consistent with the properties of STDP described above have also been carried out in humans. In these experiments, median nerve stimulation is paired with transcranial magnetic stimulation, either to change the representation in primary motor cortex, evaluated as motor evoked potentials measured electromyographi-

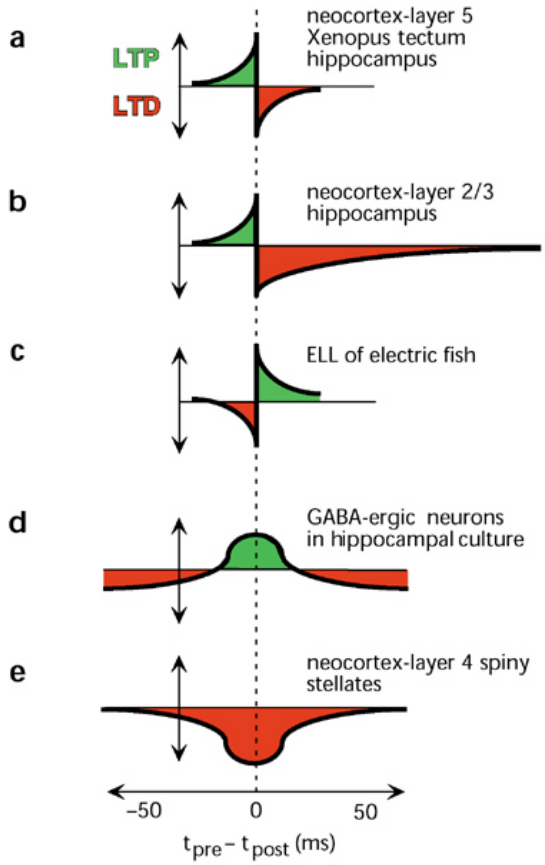

Figure 1.3. Multiple STDP Curves. Reprinted by permission from Macmillan Publishers Ltd: Nature Neuroscience (Abbott and Nelson), copyright (2000).

cally (Wolters et al., 2003), or to alter the representation in primary somatosensory cortex, measured electroencephalographically (Wolters et al., 2005).

The cellular mechanisms underlying STDP are not fully understood, and, to the extent that they are known, appear to differ somewhat depending on the synapse under consideration. According to the most conventional model, *induction* of both LTP and LTD is dependent on calcium influx through NMDARs, and subsequent activation of multiple signal transduction cascades. Most prominent among the signaling molecules involved is calcium/calmodulin-dependent protein kinase II (CaMKII), but very many others have been implicated in at least a modulating role (Lisman, 1989; Malenka and Bear, 2004). The resultant *expression* of LTP or LTD can

be effected presynaptically, e.g., changing neurotransmitter release (Zakharenko et al., 2001; Emptage et al., 2003; Ward et al., 2006) or postsynaptically, e.g., changing the number of AMPARs (Malenka and Nicoll, 1999; Malinow and Malenka, 2002; Song and Huganir, 2002), with either process having the potential to convert silent synapses into functional ones (Ward et al., 2006). Long term maintenance of LTP and LTD is associated with protein synthesis (Sutton and Schuman, 2006), gene transcription (Pittenger and Kandel, 2003) and structural remodeling of the synapse (Abraham and Williams, 2003; Malenka and Bear, 2004).

It is an open question whether the standard single coincidence detector model is sufficient to account for the functional form of Hebbian STDP (Dan and Poo, 2004). It prescribes that LTP results from a large calcium transient due to extensive NMDAR activation, and that LTD should occur if the rise in calcium is smaller in amplitude and more sustained, due to limited NMDAR activation and opening of voltage-gated calcium channels. The reason for this divergence is that, when the postsynaptic AP *follows* the EPSP, transmitter is still bound to the receptor when the AP gives rise to sufficient depolarization to relieve the Mg^{2+} block. When the AP *precedes* the EPSP, the AP afterdepolarization that now coincides with the time of transmitter binding to the receptor is of much smaller amplitude than the peak of the AP, and presumably less capable of dislodging the Mg^{2+} block. There is some direct evidence from 2-photon imaging experiments that calcium influx is supralinear if the postsynaptic AP follows the EPSP within a few 10s of milliseconds, and sublinear when the order is reversed (Koester and Sakmann, 1998). However, considering the Hebbian STDP curve for positive Δt (Figure 1.2), given that both LTP and LTD require a rise in calcium, one would predict that, as the delay between the EPSP and the postsynaptic AP increases, the associated calcium level should

traverse a depression-evoking window before reaching a level that would leave synaptic efficacy unchanged (Dan and Poo, 2006). Such a window of depression is not generally observed[4]. One hypothesis proposes that calcium levels arising from the EPSP itself, although insufficient to cause LTP, are already above the level that would lead to LTD, and that, if the EPSP is preceded by an AP, calcium influx is reduced because of NMDAR desensitization or AP afterhyperpolarization (Dan and Poo, 2006). There is little evidence, however, that a decrease in calcium, relative to baseline, can evoke LTD (Karmarkar et al., 2002), and experiments in which the AHP was reduced at the relevant Δt did not alter LTD (Feldman, 2000). Another possibility is a second coincidence detector to further distinguish the order of pre- and postsynaptic events as proposed by modeling studies (Karmarkar and Buonomano, 2002), and supported by electrophysiological (Bender et al., 2006) and imaging data (Nevian and Sakmann, 2006). In this a scenario, LTD occurs when voltage-gated calcium channels are opened by the AP, raising calcium levels by the time transmitter binds the NMDARs (and/or metabotropic glutamate receptors). This combination of dynamics and sources of calcium is generated only for $\Delta t < 0$, permitting the unique mapping onto depression (Karmarkar et al., 2002).

Modeling studies have addressed several aspects of STDP, ranging from biophysical mechanisms to algorithmic implications. A number of studies model the calcium dynamics resulting from NMDAR opening (Kitajima and Hara, 2000; Karmarkar et al., 2002; Shouval et al., 2002; Rubin et al., 2005; Kubota and Kitajima, 2007), and demonstrate that these are sufficient to account for bidirectional plasticity, although Karmarkar et al. (2002) predict some level of depression for $\Delta t > 0$, not generally observed. Several solutions have been proposed to resolve

[4]With few exceptions, e.g., Nishiyama et al., 2000

this issue, including a second Δt detector (Karmarkar and Buonomano, 2002, as discussed in some detail above); the stochastic properties of synaptic transmission itself (Shouval and Kalantzis, 2005); or an explicitly joint dependence on the peak amplitude and duration of calcium influx(Kubota and Kitajima, 2007). Rubin et al. (2005) include in their model an explicit representation of plausible biochemical cascades that serve to read out the calcium signal and convert it into potentiation or depression. They observe that, although it is possible to generate the Hebbian curve solely on the basis of the calcium signal, entails an inevitably steep dependence on the particulars of the backpropagating AP and synaptic transmission. (They suggest that this could also account for the variation in STDP observed under different experimental conditions). Kubota and Kitajima (2007) evaluate the effect of switching NMDAR subunit types (NR2A vs NR2B, characterized by different decay rates) as observed in the course of neural development, and conclude that it can account for the potentiation and stabilization of weak synapses following competition between them during a critical period. A number of experimental observations are also captured by the phenomenological model of Abarbanel et al. (2002), including synaptic strength changes resulting from periodic and Poisson distributed spike trains, as well as Anti-Hebbian plasticity.

Another set of studies analytically derives the functional form of the STDP curve by a number of different approaches, including optimization of the likelihood of postsynaptic firing at specified times by gradient ascent (Pfister et al., 2006); a reinforcement learning algorithm (Florian, 2007), modified from Baxter et al. (2001); and an objective function minimizing postsynaptic response variability in Gerstner's (2001) biophysically realistic spike-response model (Bohte and Mozer, 2007).

A large number of modeling studies are devoted also to issues regarding what is computed by networks implementing STDP. Several reports, representing a wide range of methods, describe competition among inputs giving rise to selectivity for a set of predictive and temporal coding-related properties of the networks' inputs: timing precision, earliest to arrive in a sequence, synchrony, and correlation among inputs (Song et al., 2000; Kistler and van Hemmen, 2000; Kistler, 2002; Bofill-i Petit and Murray, 2004; Guyonneau et al., 2005; Zou and Destexhe, 2007). Additional properties concomitant with the development of such selectivity include decreasing latency of postsynaptic response (Guyonneau et al., 2005), stable synaptic weight distibutions (Song et al., 2000), the stabilization of output rates (Kempter et al., 2001), synchronization of spike-volleys propagating through the network (Suri and Sejnowski, 2002), decoupling of synchronously stimulated subset of neurons from the rest of the recurrently connected balanced network (Morrison et al., 2007), and self-organized switching between rate and temporal coding (Masuda and Aihara, 2004; Masuda and Aihara, 2007). Lazar et al. (2007) report a simple recurrent network that implements both STDP and intrinsic plasticity, and instantiates time series prediction and also exhibits memory fading.

1.2.2 The Mushroom Body

Insect learning and memory have been studied predominantly in two species: the fruitfly, *Drosophila melanogaster* and the honeybee, *Apis mellifera*, from quite different perspectives. Experiments addressing learning in fruitflies have been mostly genetic, while honeybees have been studied from a behavioral (as well as neuroanatomical and limited electrophysiological) perspective. Behavior has also been assessed in *Drosophila*, but mostly in simple conditioning assays, at a population

level, to screen for mutants to identify genes or implicate subpopulations of cells. In contrast, honeybee behavior is studied in considerable detail in individual animals in the laboratory, in the hive, and in free flight. Behaviors such as conditioned proboscis extension, waggle dance, navigation and other aspects of foraging behavior have been investigated for many decades. This marked division of labor is further exemplified in the locust, *Schistocerca americana* (as well as the blowfly, *Calliphora vicina* and cockroach *Periplaneta americana*, among others), which, because of their nervous systems' robustness and accessibility, have allowed for detailed examination of neural coding issues in sensory and motor systems. Only relatively recently have these different approaches been combined in the same species, particularly in *Drosophila* (Wilson and Mainen, 2006) and to a lesser extent in the honeybee (Sattelle and Buckingham, 2006).

The structure identified for its involvement in learning and memory, more than any other region of the insect brain, is the Mushroom Body. It is a bilaterally symmetrical structure consisting of a calycal neuropil (in some species there are two per hemisphere), which forms a cup beneath a large number of Kenyon Cell somata (see Table 1.1), and a pedunculus that terminates in two or more lobes. The KCs (the MB's intrinsic cells) send dendrites into the calyx and their axons make up the pedunculus and subsequently bifurcate into the lobes (Figure 1.1B). The input to the calyx is predominantly olfactory, from Antennal Lobe Projection Neurons. In Hymenoptera, there is evidence of direct calycal input from the optic lobes as well (Jawlowski, 1958; Jawlowski, 1960); in Orthopterans, gustatory inputs from the lobus glomerulatus have also been shown to supply the calyx directly (Weiss, 1981; Strausfeld et al., 1998). Beyond the calyx, the MB also receives input in the pedunculus and in the lobes (by virtue of synapses onto KC axons, some of which

are made by other KCs). The output of the MB appears to be restricted to the lobes, where KC axons contact MB extrinsic neurons. These neurons form (input and output) contacts in other regions of the brain, and also synapse recurrently in all areas of the MB, thus providing further multimodal input. The KC population can be divided into multiple types, as determined by anatomical methods (Strausfeld et al., 1998; Farivar, 2005; Fahrbach, 2006). In the pedunculus and lobes, KC axons appear to segregate into multiple concentric or parallel layers, subsets of which are selectively invaded by individual MB extrinsic neurons. This segregation also has a counterpart in the calyx, where a KC's pattern of dendritic invasion is predictable from its axonal layer in the pedunculus and lobes In some species, different KC types are also distinguishable by the subset of lobes their axons project to.

The MB was first described by Felix Dujardin in 1850; his conception of its function, although not involving learning and memory, still bears notable similarity to the model currently held: he considered the MBs as bestowing the insect with some level of control over instinctive behaviors (Dujardin, 1850; Dujardin, 1853). His findings supporting this assertion include the result that MB size (both across and within species) appears to be correlated with the extent of social behavior in intact animals, and that decapitation affects the ability to maintain coordinated motor actions to a greater extent in animals with larger MBs (Dujardin, 1850; Strausfeld et al., 1998). A comparable view was held by F.C. Kenyon (1896), namely of the MB as an area for multi-sensory integration, quite distinct from the *direct* sensory-motor relays found elsewhere in the brain and other ganglia (Strausfeld et al., 1998). A role for the MB in associative memory was proposed later, based on data from conditioning experiments in the honeybee (Menzel et al., 1974) and *Drosophila* (Quinn et al., 1974; Heisenberg, 1980). Anatomical manip-

ulations were initially made by lesioning tissue with hydroxyurea, while later studies have employed more sophisticated genetic techniques to inactivate different aspects of the system (Heisenberg, 2003). Such tools have been used to evaluate the role of particular molecules, (e.g., proteins in the cyclic AMP signaling pathway such as adenylyl cyclase, cAMP dependent protein kinase, cAMP phospodiesterase and cAMP-response element binding protein, among others), in learning and memory. In addition, particular molecules known to be necessary for, or detrimental to neurotransmission have been employed to transiently inactivate subpopulations of cells, using the GAL4-UAS system to promote specificity and temperature sensitivity to effect transience (McGuire et al., 2005; Davis, 2005; Liu and Davis, 2006). While such studies constitute an extraordinary improvement over relatively crude early manipulations, and hold considerable promise, sufficient specificity and extent of expression are not so commonly achieved.

Chapter 2

Oscillations and Sparsening of Odor Representations in the Mushroom Body

In the insect olfactory system, oscillatory synchronization is functionally relevant and reflects the coherent activation of dynamic neural assemblies. We examined the role of such oscillatory synchronization in information transfer between networks in this system. The antennal lobe is the obligatory relay for olfactory afferent signals and generates oscillatory output. The mushroom body is responsible for formation and retrieval of olfactory and other memories. The format of odor representations differs significantly across these structures. Whereas representations are dense, dynamic, and seemingly redundant in the antennal lobe, they are sparse and carried by more selective neurons in the mushroom body. This transformation relies on a combination of oscillatory dynamics and intrinsic and circuit properties that act together to selectively filter and synthesize the output from the antennal lobe. These results provide direct support for the functional relevance of correlation codes and shed some light on the role of oscillatory synchronization in sensory networks.

2.1 Introduction

Electroencephalogram and local field potential (LFP) oscillations generally indicate periodic coherent synchronization of neuronal assemblies (Adrian, 1942; Gelperin and Tank, 1990; Gray, 1994; Bragin et al., 1995; Steriade et al., 1996; Csibra et al., 2000). While the occurrence of macroscopic oscillations has now been correlated with various sensory, behavioral or cognitive states in mammals (Gray et al., 1989; Eckhorn et al., 1988; Rodriguez et al., 1999; Patel and Balaban, 2000; Fries et al., 2000; Engel et al., 2001), the functional significance of such observations is debated (Shadlen and Movshon, 1999; Abbott and Dayan, 1999). Many hypotheses based on temporal correlations have been proposed (Abbott and Dayan, 1999; Shadlen and Movshon, 1999; Hopfield, 1995; von der Malsburg and Schneider, 1986; Diesmann et al., 1999): among others, that cortical neurons might act as coincidence detectors, rather than integrators, and thus select for correlated input (Abeles, 1982; Konig et al., 1996). Most hypotheses, however, remain tentative for lack of direct experimental test. The olfactory nervous system, in which molecular design (Clyne et al., 1999; Gao et al., 2000; Vosshall et al., 2000; Mombaerts et al., 1996), circuit architecture (Mombaerts et al., 1996) and oscillatory dynamics (Adrian, 1942; Gelperin and Tank, 1990; Laurent and Naraghi, 1994; Lam et al., 2000) appear common across phyla, offers a rare opportunity to study some of these coding issues.

2.2 Results

2.2.1 Olfactory circuits

The insect antennal lobe (AL) is the analog of the vertebrate olfactory bulb. In locusts, each AL receives input from ~90,000 ORNs and contains ~1130 densely

interconnected neurons (300 local inhibitory neurons, LNs; 830 excitatory, multiglomerular projection neurons, PNs) (Leitch and Laurent, 1996; MacLeod and Laurent, 1996). Each AL sends distributed projections to the ipsilateral mushroom body (MB), a memory area (Heisenberg et al., 1985; Dubnau et al., 2001; McGuire et al., 2001; Zars et al., 2000). PNs are the only channel for olfactory input to the MB. Conversely, there is no evidence for feedback from the MB to the AL. Each locust MB contains ~50,000 small neurons (Kenyon cells, KCs) (Laurent and Naraghi, 1994; Kenyon, 1896), whose spiny dendrites receive direct input from PNs (Laurent and Naraghi, 1994). In locusts, each PN contacts about 600 KCs (~30 synaptic varicosities per PN axon, times ~20 distinct synaptic contacts with different KC profiles per varicosity, Leitch and Laurent, 1996). Each KC receives contacts from many PNs, as seen from incremental electrical stimulation of PNs. The total number of outputs made by all PNs onto KCs must equal the total number of PN inputs received by all of KCs: Hence, if 830 PNs project to 50,000 KCs with 1:600 divergence, the average PN-to-KC convergence is on the order of 10. While unknown, these ratios probably vary by little more than a few fold across the PN/KC populations.[1] The dendritic tree of a typical KC contains 100 to 200 spines (Laurent and Naraghi, 1994). Many of these inputs must thus originate outside of the AL.

Odor-evoked PN responses exhibit globally coherent 20–30 Hz oscillations and stimulus- and PN-specific slow modulation of firing rate, both shaped in great part by LN-mediated inhibition (MacLeod et al., 1998; Stopfer et al., 1997; Wehr and Laurent, 1996; MacLeod and Laurent, 1996; Laurent et al., 1996). Hence, during a stimulus, the AL output consists of barrages of spikes from an evolving PN assembly. While individual PN spike timing during one oscillation cycle can be

[1] Recent results by Jortner et al. (2007) reveal that the PN-to-KC convergence is closer to ~400.

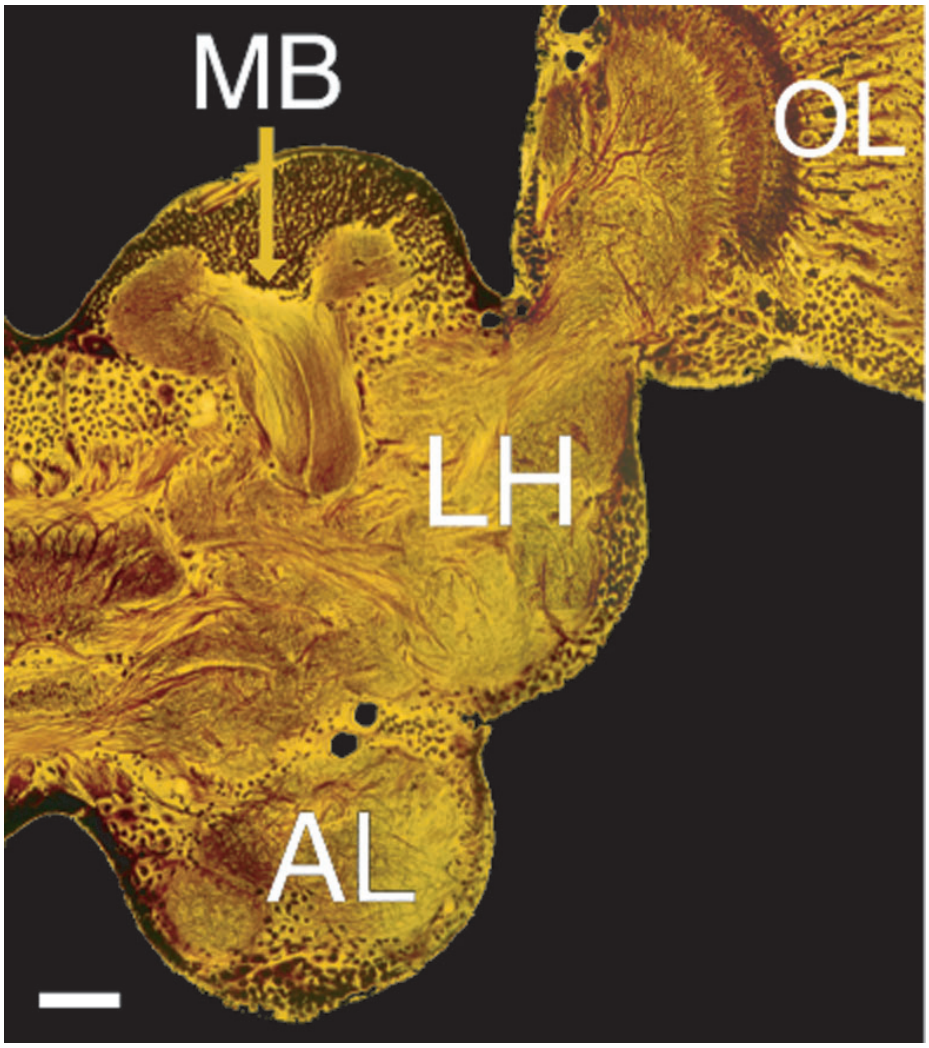

Figure 2.1. Olfactory circuits: Transverse section of the locust brain (left half, Bodian stain). Olfactory input originates from olfactory receptor neurons (ORNs) on the antenna. ORN axons terminate in the antennal lobe (AL), where projection neurons (PNs) act as relays, with projections to the mushroom body (MB) and the lateral horn (LH). OL: optic lobes. Calibration: 80 μm.

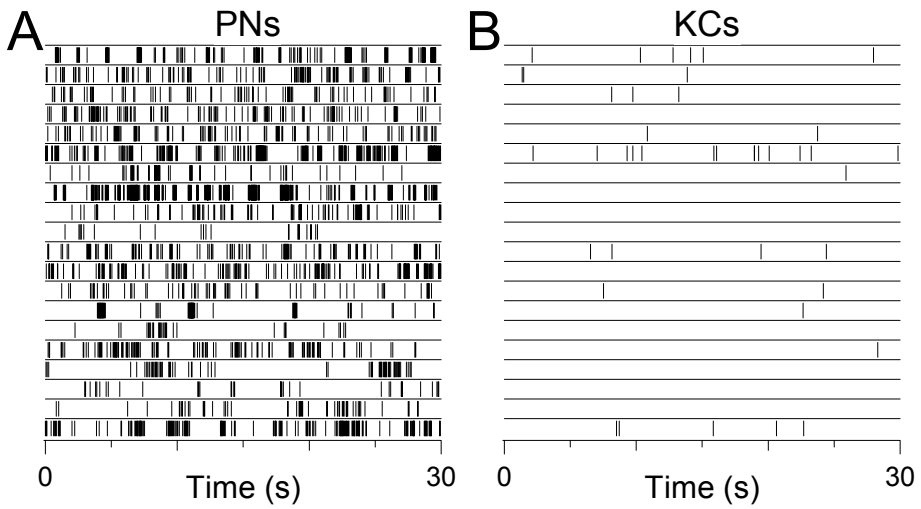

Figure 2.2. PN and KC baseline firing in the absence of odor stimulation (see text). Thirty-second rasters of 20 PNs (*A*) and 20 KCs (*B*) recorded with tetrodes. Note the exceedingly low baseline activity of KCs. (Empty rasters denote absence of action potentials during the randomly selected segment chosen for display. These rasters, however, of course originate from identified KCs, whose action potentials occurred at other times during the recording period.) [PN data (*A*): O. Mazor. KC data (*B*): J Perez-Orive.]

phase locked, it is not so for all PNs active during that cycle. At what time(s) a PN locks to others depends on both the odor and the PN. To understand the decoding of PN output by KCs, we examine the firing behavior of both populations at rest and in response to odors.

2.2.2 Resting activity

Baseline activity profiles of PNs and KCs were measured over several-minute long stretches of uninterrupted recording in naïve animals, using multiple tetrode recordings (see section 2.4.2, p. 45). At rest, the PN population fired at a mean rate of 3.87 ± 2.23 spikes/s per PN (range: 0.49 to 10.4, n = 35 PNs). Baseline firing was >100 times lower in KCs (median: 0.025 spike/s, interquartile range: 0.088 spike/s; n = 23 KCs) (figure 2.2). Hence, despite a constant excitatory drive from PNs, KCs at rest remained remarkably inactive.

2.2.3 *Response selectivity*

PNs and KCs were challenged in awake animals with a panel of odors (typically 17; range: 5–24; 5–25 trials per odor; 1 s pulses; 20–30 s between trials; see section 2.4.1, p. 44). Experimental conditions were identical for PN and KC recordings.

SPIKING RESPONSE PROBABILITY

The probability of observing a stimulus-evoked change in firing behavior was different across the two populations (figure 2.3). Most PNs exhibited a reliable change in firing behavior within the first few seconds following stimulus onset. They showed complex temporal patterning (with increases and decreases in instantaneous firing rate) that often greatly outlasted the stimulus itself (figure 2.3A). Many of these responses were inhibitory, and many of these inhibitory periods were followed by a period of increased firing, up to five seconds after stimulus offset. We analyzed excitatory response probabilities across PNs (and KCs) quantitatively, using a variety of methods and analysis windows. We show here the results obtained with Method A (see section 2.4.4, p. 48). Results obtained with the other methods are nearly identical (table 2.1 and figure 2.10). The distribution of response probabilities for PNs was broad (figure 2.4A), with means over all cells of 0.64 + 0.32 (n = 58 PNs, 1140 PN-odor pairs). KC responses to these same odors were extremely rare: over all KCs (n = 74 KCs, 1101 KC-odor pairs), 58% failed to show any detectable response to any of the odors presented (figure 2.4A). The distribution of response probabilities was heavily skewed towards low values (figure 2.4A), even when considering only those KCs that produced at least one response. The mean response probability was 0.11 after averaging all KCs' individual response probabilities (figure 2.4A) (median: 0.00; interquart. range 0.12). Figure 2.3B shows three typical responsive

KCs. Among all recorded KCs, only two responded to all odors presented (10 and 12 odors, respectively). To avoid possible sampling bias, recordings were made from all regions and depths of the KC soma layers. Responsive and unresponsive KCs were found everywhere, consistent with the anatomy of PN axonal projections in the MB (Laurent and Naraghi, 1994). Similarly, no selection bias towards strong responses existed, for the great majority of them were extremely brief (c.f., below) and rarely detected on line. Selective and promiscuous KCs could occur simultaneously on the same tetrode, indicating that differences in tuning width were not caused by global modulation of excitability over time.

RESPONSE INTENSITY

Response patterns and intensities differed in PNs and KCs. While PN responses often lasted several seconds (figure 2.3A), KC responses were brief and lacked the slow temporal patterning typical of PNs (figure 2.3B). Using responsive cell-odor pairs, we counted action potentials produced by PNs and KCs over the 3 s window after stimulus onset. The distribution of PN spike counts over that period was broad, with a mean of 19.53 ± 10.67 spikes. KCs responded with 2.32 ± 2.68

Figure 2.3 ***(on the next page).*** Tetrode recordings of odor responses in PNs (*A and C*) and KCs (*B, D, and E*). In all panels, shaded area = odor puff = 1 s. *A*, responses of three simultaneously recorded PNs (PN1–PN3) to 16 different odors (first 10 trials with each stimulus displayed). Odors from top, left column: hpo, don, che, hx3, unn, min, oca, pnn; right column: chx, oco, nnn, thx, 2hp, nna, 3hp, hxo; abbreviations in methods, p. 44). *B*, responses of three KCs to the same 16 odors. Conditions as in A except: for six of the odors, KC1 and KC2 have only five trials; in KC2, the 7th odor in the right column is hxa. *C*, expanded view of PN1 raster in response to hxo (trials 3 to 15). Note alignment of spikes. *D*, response of a fourth KC to hx3 (trials 3 to 15). Note low baseline activity and alignment of first spike in the response across trials. *E*, response of a fifth KC with superimposed LFP, recorded in the MB (10–55 Hz bandpass). Note phase-locking of KC spikes. LFP: 200 μV. [PN data (*A, C*): O. Mazor and S. Cassenaer. KC data (*B, D, E*): J Perez-Orive.]

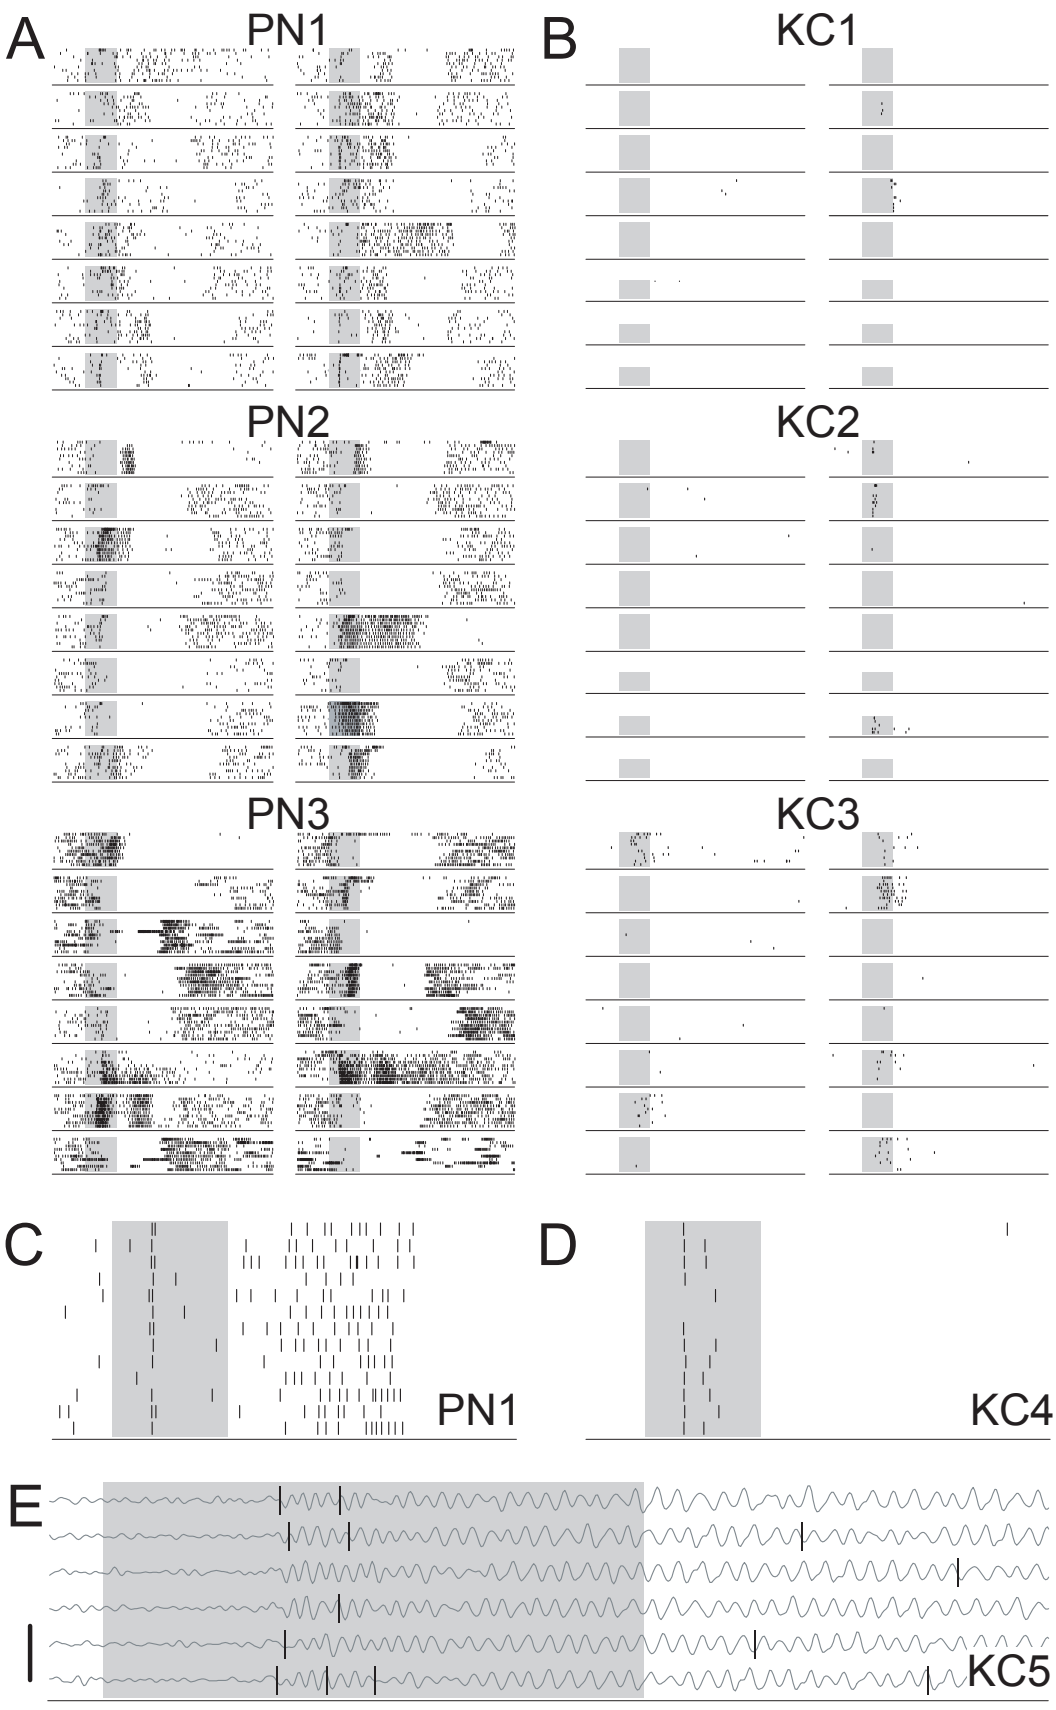
A
PN1
PN2
PN3
B
KC1
KC2
KC3
C
PN1
D
KC4
E
KC5

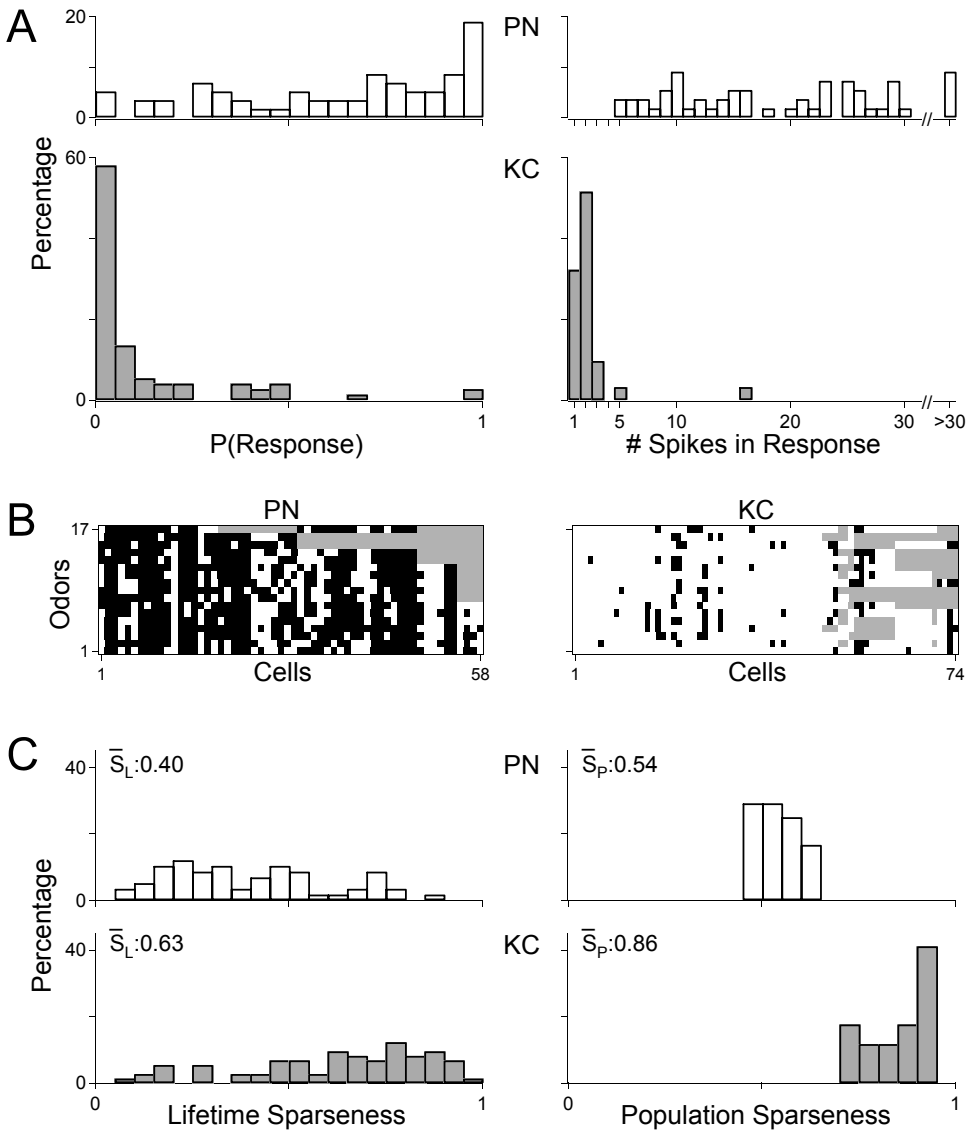

Figure 2.4. Statistics and sparseness of PN and KC odor responses (see methods). *A left*, probability of responding to x% of all odors tested (x in 5% bins) (see section 2.4.4, p. 48); note opposite skew in PN and KC distributions. *A right*, response intensity distributions (1 spike bins, measured over 3 s window). Spike counts were computed only from cell-odor pairs with a detected excitatory response during the analysis window. *B*, excitatory responses (filled squares) of individual PNs and KCs (columns) (n = 58 PNs, 74 KCs) to 17 different odors (rows: hx3, thx, chx, hxo, hpo, oco, nna, nnn, don, pnn, 2hp, 3hp, oca, unn, che, min, hxa; abbreviations in section 2.4.1, p. 44). Open squares denote inhibition (PNs only) or absence of a response. Grey squares: not tested. *C*, distributions of lifetime (left) and population (right) sparseness, computed across all cells and all tested odors. S_L and S_P are significantly different across PNs and KCs ($p < 0.001$, t'-test for S_P, z statistic for S_L). [PN data: O. Mazor and S. Cassenaer. KC data: J Perez-Orive.]

spikes[2] (figure 2.4A). We found a negative correlation between KC spike count and response selectivity (Spearman ranked correlation coefficient: -0.567, $p < 0.05$).

TEMPORAL PRECISION

PN spike probability and precision is PN-, odor- and time-specific (MacLeod et al., 1998; Stopfer et al., 1997; Wehr and Laurent, 1996; MacLeod and Laurent, 1996; Laurent et al., 1996). Time-locked PN spikes were easily detected when they occurred in isolation (e.g., figure 2.3C), but were found also within sustained responses, consistent with previous intracellular results. In KCs, individual responses typically contained about two spikes (figure 2.4A), one of which at least could be precisely locked to stimulus onset with a fixed delay. Stimulus-locked spikes were often the first ones in the KC's response, but could occur at any cycle. The first spike in the response of KC4 (figure 2.3D), for example, had a jitter of only ±4 ms relative to stimulus onset. Stimulus-locked spikes with such small jitter, however, were not commonly observed. Another measure of precision, more relevant to this system, is the timing of each action potential relative to its LFP oscillation cycle (phase) (figure 2.3F). The mean phase of KC spikes was 83° ± 77° (n = 18 KCs; where 0° is oscillation peak, figure 2.6F). Mean spike phase was the same in the most as in the least specific KCs (90° ± 67° vs. 86° ± 81°, $n = 5$ cells each). The spikes within a doublet (or triplet) were always separated by one to a few oscillation cycles (e.g., figure 2.3F). This indicates that appropriate PN drive to individual KCs lasted several oscillation cycles, and that when a KC spike was fired, it occurred preferentially at the same phase of its oscillation cycle.

[2]Most KC spikes occurred in the beginning of the response: response intensity was 2.33 ± 2.02 spikes over the first 1.4 s, while PNs produced 12.84 ± 7.29 spikes on average in that period.

2.2.4 Sparseness of odor representations across PNs and KCs

Figure 2B compresses the responses of 58 PNs and 74 KCs to the same 17 odors and illustrates the contrast between the two population representations.[3] A simple estimate of population sparseness (S_P) is the proportion of cells unresponsive to each stimulus, averaged over all stimuli. It thus represents the sparseness of the representation of each odor across the population, averaged over all odors, but ignores the strength of each response. S_P was 0.90 in KCs and 0.33 in PNs. S_P can also be calculated more directly using the neurons' firing rate distributions (Willmore and Tolhurst, 2001) for each tested stimulus, whether we detected a response or not (see section 2.4.4, p. 48). Applied to PNs and KCs, this measure of S_P was again always greater in KCs (figure 2.4C). Finally, sparseness can be calculated for each cell across all the stimuli it has experienced. This measure, called lifetime sparseness, S_L, approximates the mean tuning width of each neuron, averaged over all neurons. Again, S_L was significantly higher in KCs than in PNs (0.63 vs. 0.40, $p < 0.001$, t'-test, figure 2.4C). S_L and S_P were also calculated using the other response analysis windows, or using only the odor responsive cells. By all measures, odor representations were always significantly sparser across KCs than PNs (table 2.1 and figure 2.10).

2.2.5 Mechanisms underlying sparsening

Subthreshold KC activity during odor stimulation. Sharp electrode recordings (see section 2.4.6, p. 51) from KCs (n = 29) revealed odor-evoked, subthreshold activity made up of periodic synaptic potentials (figure 2.5A). These were locked to

[3] Responses were determined here according to Method A, (see section 2.4.4, p. 48). Nearly identical results were obtained if responses were assessed using different criteria, adapted to each population.

the LFP (figure 2.5B) and superimposed on a noisy and irregular synaptic background, away from firing threshold. Appropriate odor-KC combinations revealed reliable and time-specific EPSPs and/or action potentials. The response of the KC in figure 2.5C for example, contained a train of prominent EPSPs, late within the stimulus. One of these EPSPs led to an action potential in half of all trials with that odor. A different KC responded to the same odor with at least two reliable action potentials, at cycles 1 and 3 of the response, whether the neuron was at rest (figure 2.5D) or held depolarized by current injection. In all tested KCs, the existence, timing and reliability of these firing events were odor specific. We noted that a large component of the odor-evoked activity in KCs was inhibitory: if the KC was held depolarized by current injection, periodic hyperpolarizing potentials could be seen during a response; if the KC was held above firing threshold, odor-evoked inhibition interrupted this tonic firing (figures 2.5A and C). Odor stimulation thus also causes synaptic inhibition of KCs. Finally, the amplitude of odor-evoked EPSPs paradoxically increased when the KC was held depolarized (figure 2.5A), suggesting active membrane properties. We examined the possibility that synaptic inhibition and KC active conductances work together to make KCs coincidence detectors of PN input.

SOURCE OF MASKED ODOR-EVOKED INHIBITION

Because direct effects of PNs are excitatory and because locust PNs do not contain GABA (Leitch and Laurent, 1996), the source of odor-evoked KC inhibition should be downstream of PNs. In addition to sending collaterals into the mushroom body, PN axons terminate in the lateral horn (LH) (Hansson and Anton, 2000). We identified among their targets there a cluster of ~60 GABA-immunoreactive neurons (Lateral Horn Interneurons, LHIs), with direct axonal projections to the MB (see

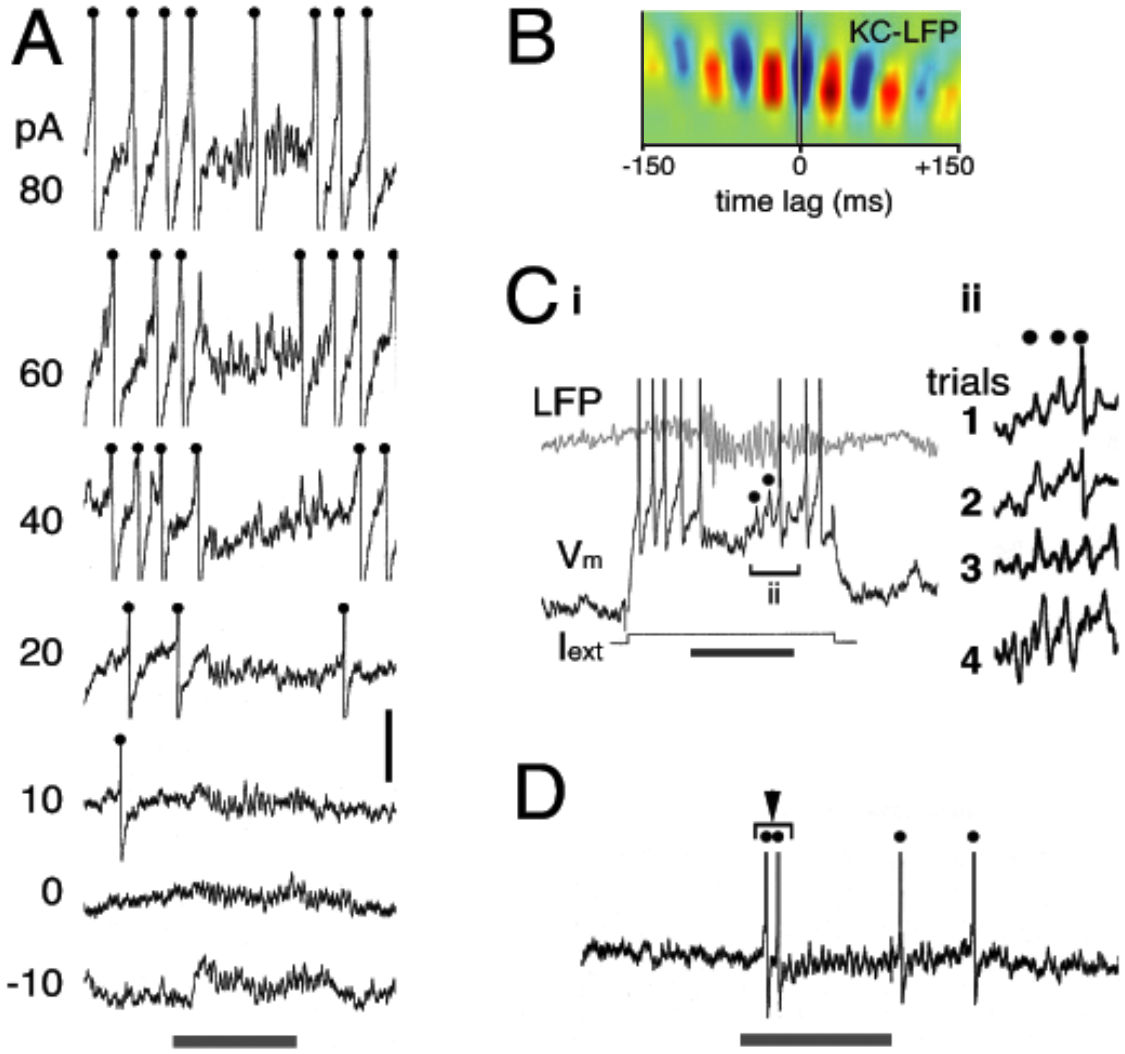

Figure 2.5. In vivo sharp-electrode intracellular records from different KCs. All action potentials are clipped. *A*, responses obtained while resting voltage set by holding currents between −10 and +80 pA. Horizontal bar: odor (cherry) delivery, 800 ms. This KC never produced any action potential in response to this odor at resting potential. Note oscillating membrane potential at rest (0 pA), interruption of DC-evoked firing by odor delivery (+20 to 80 pA traces), amplification of many discrete depolarizing potentials at most depolarized holding potentials. *B*, sliding cross-correlation of KC V_m and simultaneous LFP (different KC from *A*). Red: maxima; blue: minima. Y-axis: time, 0.5 s. Note locking of signals during odor puff. *C*, third KC recording, showing interruption of current-evoked firing by odor response and prominent, late EPSPs (•); 800 ms odor delivery (*i*). Repeated trials (1–4) show precise re-occurrence of these EPSPs during same epoch of the response; time calibration: 500 ms (*ii*). *D*, fourth KC and its spiking response to cherry odor at cycles 1 and 3. 800 ms odor pulse. Vertical calibrations: KC: 10 mV (*A, D*), 8 mV (*Ci*), 6 mV (*Cii*); LFP: 300 μV, 1–40 Hz bandpass (*C*). [KC data (*A, C, D*): G. Laurent. Cross-correlogram (*B*): G. Turner.]

section 2.4.7, p. 52) (figure 2.6A). Intracellular staining of individual LHIs showed profuse axonal collaterals, overlapping with KC dendrites (figure 2.6B). KC dendrites receive GABAergic input (Leitch and Laurent, 1996). LHIs thus appeared well suited to be a source of the odor-evoked inhibitory inputs.

LHI RESPONSES TO ODORS

LHIs responded vigorously and reliably to odors (Fig. 4C,D). LHI membrane potential oscillated in phase with the LFP (figure 2.6E), and when sufficiently excited, LHIs fired one or a short burst of action potentials at each oscillation cycle (figure 2.6C). In each cycle, LHI mean firing-time lagged 173° behind that of PNs (figure 2.6F). LHI firing phase was independent of odor identity. Synaptic drive

---

Figure 2.6 ***(on the next page).*** Feed-forward inhibition of KCs by LHIs. *A*, anti-GABA immuno-labeling (see section 2.4.7, p. 52). Cluster of ~60 reactive somata (LHI) and tract of LHI axons running to the MB (stipples). The terminals of one of these axons in the MB are shown in *B*. Calibration: 100 μm. *B*, PN axon (black) projects to the mushroom body calyx (orange) (Laurent and Naraghi, 1994) and to the lateral horn (LH). LHI (green) project to the calyx (this study). PN and LHI axons terminate on KC dendrites (red). Neurons stained by iontophoresis of cobalt hexamine (KC, PN) or neurobiotin (LHI), in separate preparations and drawn using a camera lucida. Note varicosities in LHI and PN axon collaterals. (*): KC axon. Calibration: 50 μm. *C*, representative odor-evoked responses of two LHIs and simultaneously recorded LFPs (5–40 Hz bandpass). Note membrane potential oscillations, locked to the LFP. Identity and delivery (1 s long) of stimulus indicated by black bar. Calibration: (LHI) 20 mV; (LFP) 400 μV; 200 ms. *D*, instantaneous firing rate of LHI1 (in *C*) in response to various odors. Lower edge of profile: mean instantaneous rate, averaged across trials; profile thickness: SD. All LHIs responded to all odors tested, with response profiles that varied little across different odors. *E*, sliding cross-correlation between LFP and LHI2 traces (spikes clipped). High correlation values in hot colors, low in cold. Strong locking is present throughout the response (odor delivery: vertical bar). Lower edge of correlation stripes just precedes stimulus onset due to width of correlation window (200 ms). *F*, phase relationships between PN, KC and LHI action potentials, and LFP. Polar plots: LFP cycle maxima defined as 0 rad, minima as π rad (PNs: 3 cell-odor pairs, 388 spikes; LHIs: 17 cell-odor pairs, 2632 spikes: KCs; 18 cells, 862 spikes). Mean phases shown in red. Gridlines are scaled in intervals of 0.10 (probability per bin). Below: diagram showing LFP and mean firing phases, Q. *G*, circuit diagram. [LHI data (*A–F*): G. Turner. PN data (*F*): M. Westman. KC data (*F*): J. Perez-Orive.]

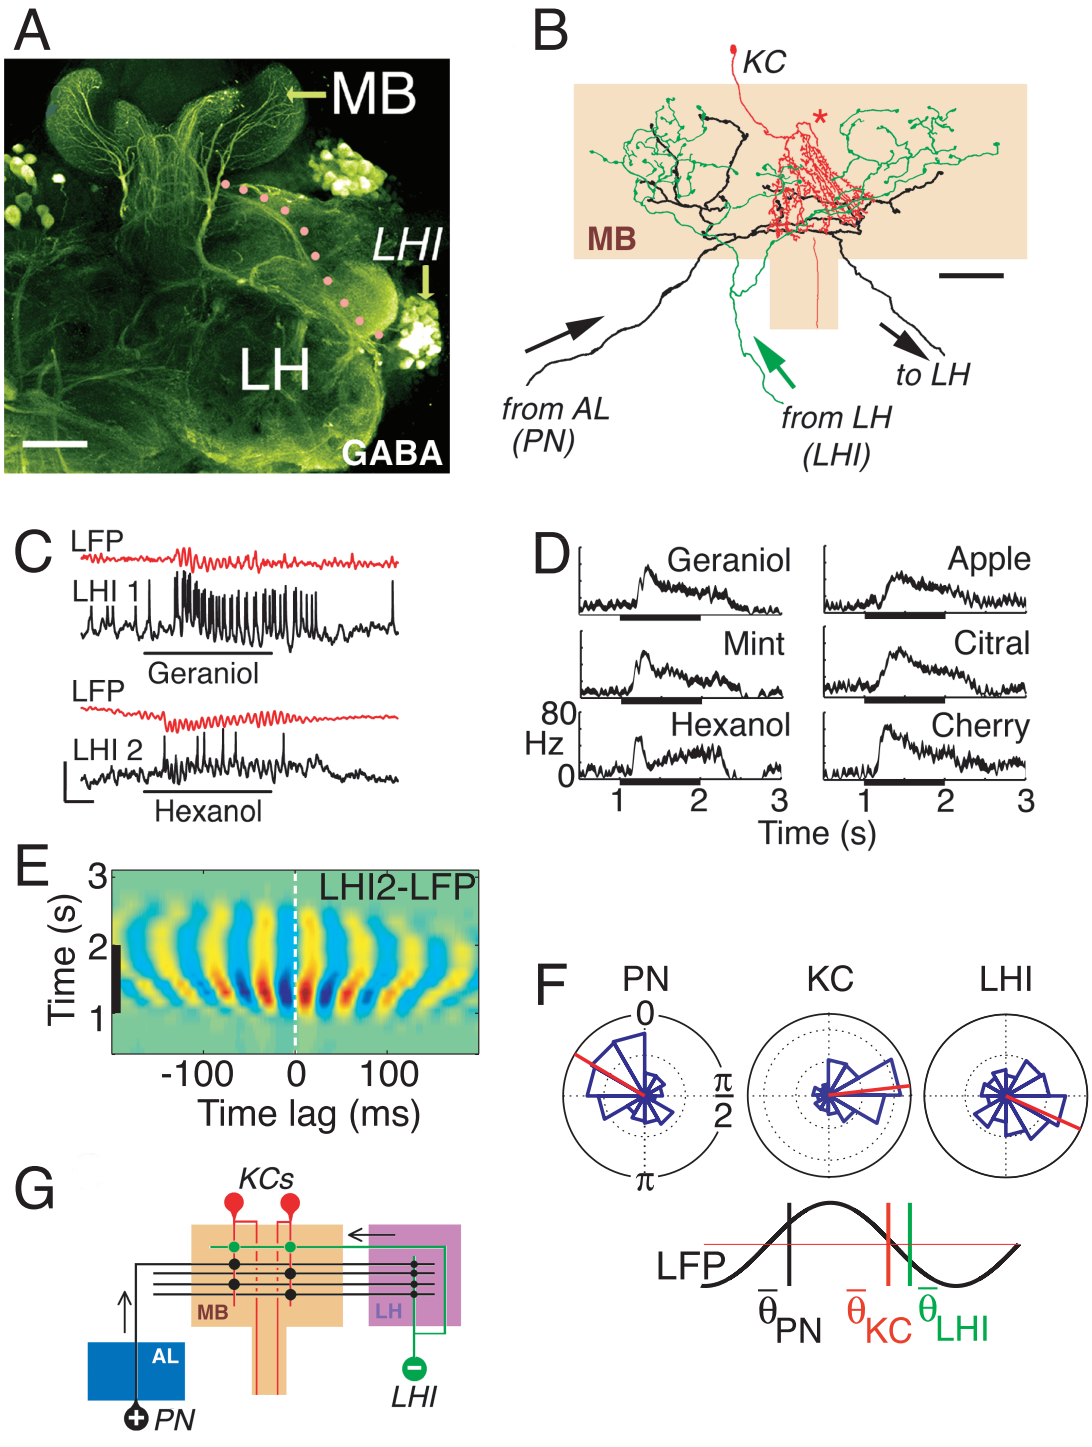
A
MB
LHI
LH
GABA
B
KC
MB
from AL
(PN)
from LH
(LHI)
to LH
C
LFP
LHI 1
Geraniol
LFP
LHI 2
Hexanol
D
Geraniol
Apple
Mint
Citral
Hexanol
Cherry
80
Hz
0
1
2
3
Time (s)
E
LHI2-LFP
Time (s)
3
2
1
-100
0
100
Time lag (ms)
F
PN
KC
LHI
0
π/2
π
LFP
θ̄PN
θ̄KC
θ̄LHI
G
KCs
MB
LH
AL
PN
LHI

to KCs thus likely consists of EPSPs from PNs alternating with IPSPs from LHIs, occurring preferentially in opposite halves of each oscillation cycle. PN and LHI inputs to KCs differ in one important respect: because each KC on average receives inputs from a very small fraction of the PNs and because the firing probability and phase-locking of each PN typically evolves during a response, the probability that many of the PNs presynaptic to a given KC fire together within the same half of one oscillation cycle is low. By contrast, individual LHIs showed sustained responses to all odors presented (figure 2.6D), consistent with the fact that 830 PNs converge onto only ~60 LHIs. Because LHI axons diverge profusely in the mushroom body (figure 2.6B), individual KCs should receive periodic input composed of consistent IPSPs, alternating with EPSPs whose total strength depends strongly on the stimulus.

SHARPENING OF KC RESPONSE TO DIRECT PN STIMULATION

We next tested more directly whether both synaptic inhibition and intrinsic active conductances assist coincidence detection in KCs. To study single EPSP-IPSP cycles in isolation, we used direct electrical stimulation of PNs rather than odors. Evoked postsynaptic potentials in KCs dramatically changed shape and duration upon varying stimulus strength (figure 2.7A). At high stimulus intensities, a sharp *spikelet* rode atop the depolarizing potential, suggesting active conductances (figure 2.7A, top trace). This spikelet was not an artifact of unusually strong stimuli: when a weak stimulus was used to elicit a smaller EPSP and holding current was adjusted so that the KC was near firing threshold, spikelets could also be observed[4] (figure 2.7B) (see section 2.4.8, p. 53; Schafer et al., 1994). Next, we tested

[4] Although we have not characterized this spikelet pharmacologically, its shape and all-or-none waveform suggest the involvement of voltage-dependent conductances (possibly Na^+ or Ca^{2+} for

the idea that GABAergic feed-forward inhibition also shapes PN-evoked PSPs. At voltages below spikelet threshold, EPSP shape was still strongly voltage dependent (figure 2.7C). Local injection of picrotoxin (PCT), a $GABA_A$-like chloride channel blocker, into the mushroom body calyx (see section 2.4.9, p. 53) broadened the EPSP and decreased the voltage-dependence of EPSP shape (figure 2.7C). This indicates that the LHI-mediated IPSP normally contributes partly, but maybe not entirely, to the shape and duration of PN-evoked EPSPs. Blocking inhibition in the calyx increased the scatter of KC-spike times following PN stimulation (figure 2.7D). LHI-mediated IPSPs thus contribute to shortening the epoch during which a KC remains depolarized after each volley of PN excitation; it could also explain why KC action potentials are so precisely phase-locked during responses to odors (figure 2.6F). Hence, the tendency of each KC to convert its excitatory input from PNs into an action potential can be facilitated in the early phase of the compound EPSP by voltage-dependent depolarizing nonlinearities and antagonized shortly thereafter by feed-forward inhibition. The remaining voltage dependence of the EPSP after PCT injection (figure 2.7C) suggests the existence of an active repolarizing conductance. Active and synaptic properties thus both likely contribute to making KCs prefer coincident input, on a cycle-by-cycle basis.

2.2.6 Influence of feed-forward inhibition on KC responses to odors

If feed-forward inhibition competes with and resets the periodic excitation of KCs by PNs, antagonizing LHI-mediated inhibition should decrease KC specificity to odors. KCs recorded in vivo with tetrodes were tested with up to 17 odors (ten

depolarization and K+ for repolarization), consistent with previous patch-clamp studies in vitro (Schafer et al., 1994)

trials per odor) and re-tested immediately after PCT injection into the mushroom body (figure 2.8). PCT caused no significant change in the KC baseline-firing rate (medians: 0.018 spikes/sec after PCT vs. 0.005 before, n = 12 KCs, p = 0.19, non-parametric sign test). PCT caused a broadening of KC tuning, characterized by greatly reduced odor selectivity (figure 2.8A–C). Even in KCs that responded to none of the odors presented in controls, responses to these same odors appeared after PCT (figure 2.8A–C). Individual KCs did not become responsive to all odors, but rather, to a larger subset of all tested odors. The mean population and lifetime

---

Figure 2.7 ***(on the next page).*** KC responses to electrical stimulation of PNs. *A*, PNs were stimulated directly using an electrode placed in the AL and evoked EPSPs were recorded intracellularly from KCs. Three traces show EPSPs recorded at progressively stronger stimulus intensities (bottom to top). Note positive inflexion during rising phase of the top EPSP and sharp repolarization. Bar graph compares EPSP half-width at the maximum stimulus intensity that was still below action potential threshold vs. half-width at 70% (± 5%) of this maximum intensity. EPSP half-width was significantly different at these two stimulus intensities ($p < 0.001$, paired t-test, n = 11 KCs). *B*, intrinsic active conductance amplifies and sharpens EPSPs near threshold. KC held near threshold with constant holding current; constant PN stimulus amplitude: successive trials elicited full-blown sodium spikes (light grey), subthreshold EPSPs (black), or intermediate *spikelet*. Sample traces collected in picrotoxin; similar spikelets were observed in control conditions (e.g., *A*). *C*, synaptic inhibition shortens KC EPSP. At progressively depolarized holding potentials, EPSP half-width significantly decreased (half-width at −40 mV significantly smaller than half-width at −60 mV, $p < 0.0005$, paired t-test, n = 10 KCs); all analyzed data below threshold for spikelet activation. After picrotoxin injection in MB, EPSPs became broader (−60 mV half-width significantly increased in picrotoxin, $p < 0.05$, t-test, n = 9). EPSP shape was less dependent on postsynaptic voltage (−40 mV half-width as percentage of −60 mV half-width significantly increased in picrotoxin, $p < 0.05$, t-test, n = 9), but was still voltage dependent ($p < 0.05$, t-test, n = 9). Sample traces from two KCs in same brain. *D*, synaptic inhibition narrows the window in which KCs can fire following PN stimulation. Stimulus intensity was adjusted to elicit an EPSP of 5–10 mV (when KC held at −60 mV), and then holding current was adjusted so that this EPSP elicited a spike on 30–60% of trials. Representative traces (left) show those sweeps that elicited spikes (arrows mark stimulus, bars mark interquartile range of spike times encompassing the difference between the 25th and the 75th percentile). Sample traces from two KCs in same brain. Group data (right) shows the interquartile range for each cell. Picrotoxin significantly increased the magnitude of the interquartile range ($p < 0.05$, t-test, n = 6 control KCs, 6 KCs in picrotoxin). *B–D*: whole-cell recordings (see section 2.4.8, p. 53). [Sharp microelectrode recordings (*A*): J. Perez-Orive. Whole-cell patch recordings (*B–D*): R. Wilson.]

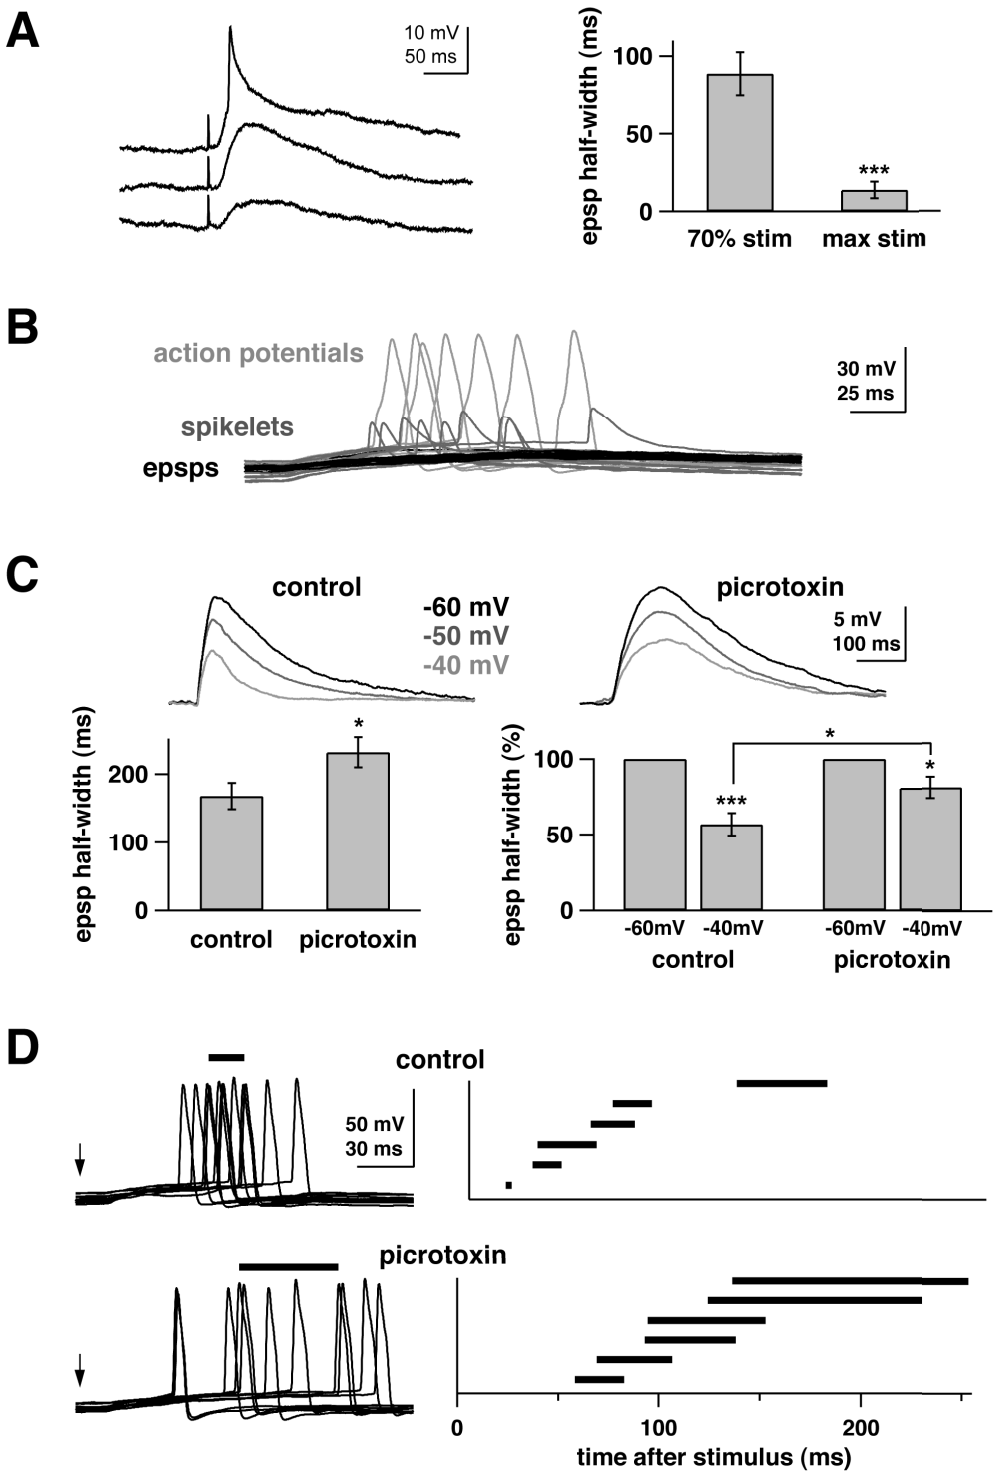
A
10 mV
50 ms
epsp half-width (ms)
100
50
0
***
70% stim
max stim
B
action potentials
spikelets
epsps
30 mV
25 ms
C
control
-60 mV
-50 mV
-40 mV
picrotoxin
5 mV
100 ms
epsp half-width (ms)
200
100
0
*
control
picrotoxin
epsp half-width (%)
100
50
0
*
*
***
-60mV
-40mV
-60mV
-40mV
control
picrotoxin
D
control
50 mV
30 ms
picrotoxin
0
100
200
time after stimulus (ms)

sparseness calculated over this KC subset was significantly decreased after PCT (S_P = 0.70 to 0.41, n = 11 odors, $p < 0.001$, paired t-test; S_L = 0.47–0.30, n = 12 KCs, $p < 0.05$, non-parametric Wilcoxon signed-rank test). Individual KC response intensity after PCT treatment was not significantly different from control (1.96 ± 0.81 spikes; PCT: 1.82 ± 0.47 spikes), but KC action potentials after PCT lost their locking to the LFP (figure 2.8D, cf. controls, figure 2.6F). This confirms earlier experiments (figure 2.7D) suggesting that LHI-mediated IPSPs normally constrain KC integration and spike timing.[5]

2.3 Discussion

In the AL, individual odors are represented by a large fraction of the 830 PNs: Baseline activity is high, sparseness is low, individual PN responses are sustained. In the MB, the same odors activate a small proportion of neurons in a larger population (50,000 KCs): baseline activity is close to zero, sparseness is high, individual KC responses are rare and typically contain two action potentials only. KC action potentials thus each carry much more information than those of PNs.

how does sparsening arise?

We propose that KCs act as selective coincidence detectors on periodic PN input: Because individual KCs receive inputs from only a small fraction of all PNs, because the patterned responses of individual PNs are staggered in time and because EPSP summation by KCs occurs best within a fraction of each oscillation cycle, the conditions appropriate for bringing a KC to threshold are rarely met. During odor

[5]PCT application to the mushroom body did not affect the LFP oscillations recorded there, for the principal source of these oscillations—synchronized, periodic synaptic input drive from PNs—was excitatory and cholinergic (nicotinic).

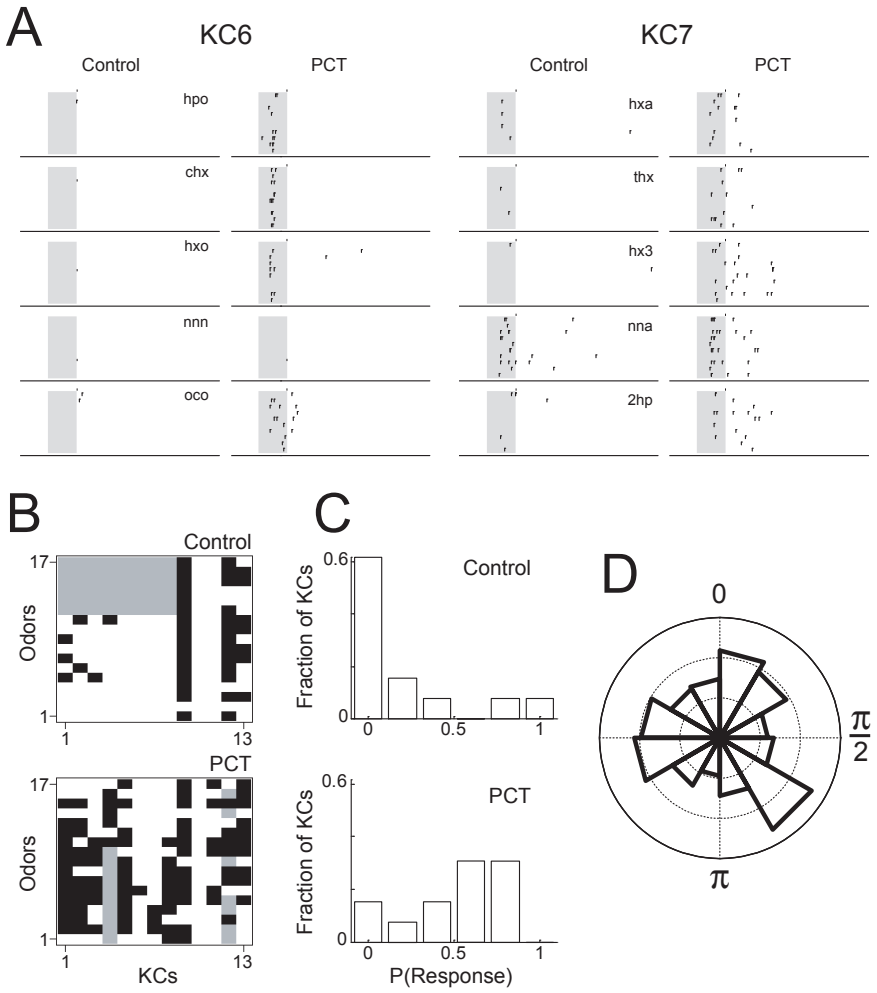

Figure 2.8. Influence of feed-forward inhibition on KC odor tuning and phase-locking (in vivo, wire tetrode recordings). *A*, two KCs (5–6) and their responses to five odors before and after local PCT injection into the MB (see section 2.4.9, p. 53). Odor pulses (shaded area): 1 s; 10 trials per odor, top to bottom (abbreviations in section 2.4.1, p. 44). *B*, comparison of KC response profiles before and after PCT. Filled: response; open: no response; grey: not tested (more odors were generally tested after PCT treatment). PCT broadened KC response tuning profiles, by causing the appearance of responses to new odors. Odors, 1 to 17: oca, hxa, thx, hx3, oco, unn, nna, 2hp, che, chx, hxo, don, nnn, 3hp, hpo, pnn, min. *C*, frequency distribution of response probabilities (across all odors tested) before and after PCT treatment (n = 13 KCs). Note dramatic reduction of proportion of specific KCs (leftmost bin) after PCT. Median response probabilities: 0.09 (control), 0.59 (PCT). *D*, phases of KC spikes relative to LFP, during odor-evoked responses (0 rad: max; π: min of LFP voltage). Vector strengths: 0.03 (PCT) vs. 0.41 (control, Fig 4F). Gridlines in intervals of 0.05 (probability per bin). [J. Perez-Orive]

stimulation, each oscillation cycle contains both locked and unlocked PN spikes (Laurent et al., 1996). Periodic IPSPs, caused in KCs by LHIs whose mean firing is in antiphase with the discharge of the synchronized PNs, antagonize the action of inappropriately timed PN action potentials. When LHI-mediated inhibition is blocked, this normally antagonized excitatory drive to KCs can now summate over a longer time window: KCs lose much of their specificity. Time-locked feed-forward inhibition thus helps define very short but renewed (once per oscillation cycle) integration windows for each KCs, akin to a periodic reset, with critical consequences for KC specificity.

None of the features uncovered so far (oscillatory patterning, feed-forward inhibition, fan-in and fan-out, active properties) are unusual ones (Fricker, 2000; Galarreta and Hestrin, 2001; Pouille and Scanziani, 2001; Contreras et al., 1997; Haberly, 1990). In particular, distributed and partly overlapping projection patterns of mitral cells have been seen in rodent prepiriform cortex (Zou et al., 2001) and local feed-forward inhibitory circuits are common (Fricker, 2000; Pouille and Scanziani, 2001; Contreras et al., 1997; Haberly, 1990). Nonlinear intrinsic properties have been seen in some cortical and hippocampal cells among others and hypothesized to underlie coincidence detection (Fricker, 2000; Galarreta and Hestrin, 2001). We showed here that all these properties exist together in the same circuit and that their concerted use in the context of oscillatory activity results in a major transformation of sensory codes.

how could sparsening be useful?

Because the mushroom body is a likely site for the formation and retrieval of olfactory memories (Heisenberg et al., 1985; Dubnau et al., 2001; McGuire et al., 2001;

Zars et al., 2000), we must ask why sparse codes might be advantageous there. While it is clear that extremely sparse codes ("grandmother" schemes, Barlow, 1969) may be undesirable because they confer sensitivity to damage and low capacity, representations carried by small subsets of neurons offer many theoretical advantages. First, overlaps between individual representations are less likely than if each representation used a large proportion of the available neurons, limiting interference between memories. This system's memory capacity can still be very high, because the total population size is large and sparseness is not extreme. Second, comparisons between stimulus-evoked patterns and stored memories are simpler if they invoke fewer elements. Third, representations become more synthetic or high-level. Every KC action potential compresses the signals carried by several PNs that are each potentially more informative about stimulus composition. Sparsened representations thus contain less explicit detail. This conclusion agrees with behavioral and psychophysical observations in humans, rats and insects that odor perception has a prevalent synthetic quality (Livermore and Laing, 1996; Linster and Smith, 1999; Cain and Potts, 1996).

SIGNIFICANCE FOR NEURAL CODING

Our results have implications for the understanding of neural codes. First, single-neuron responses can be exquisitely specific, extremely short (1–2 spikes only), and temporally precise (both within and across oscillation cycles). Whether response characteristics similar to those shown for KCs exist in other systems is not known; because they are so brief and specific, such response patterns are de facto hard to uncover. Studies of primate frontal and motor cortices that show very brief firing events, however, are consistent with some of our results (Abeles et al., 1993;

Riehle et al., 1997). Second, subtle yet highly relevant activity patterns may go undetected with many large-scale brain-activity monitoring techniques: sparse and brief activity is unlikely to be reflected in most macroscopic signals. Yet, as we show here, this may sometimes be all there is. Finally, to measure the relevant information content of an action potential, one must know how downstream targets interpret it. For example, we showed previously that projection neuron action potentials typically phase lock to the LFP only during certain (stimulus- and PN-specific) epochs of a response (Laurent et al., 1996). Our results indicate that KCs will be more sensitive to phase-locked PN action potentials than to those occurring closer to each LHI-mediated IPSP, whose timing is itself determined by the locked-PN population. PN spikes, therefore, are not all equally meaningful to a KC. Even in cases where firing rates are high, many spikes may be of minimal significance to a target, because improperly-timed. Here, relevance is determined by inter-neuronal correlation. Hence, deciphering brain codes requires an evaluation of these correlations and their consequences on the channeling of information. Conversely, macroscopic oscillations may indicate the existence of neural filters, whose properties will determine the interpretation one should make of a spike train.

2.4 Methods

2.4.1 Preparation and stimuli

Results were obtained from locusts (*Schistocerca americana*) in a crowded, established colony. Young adults of either sex were immobilized, with one or two antennae intact for olfactory stimulation. The brain was exposed, desheathed and superfused with locust saline, as previously described (Laurent and Naraghi, 1994).

Odors were delivered by injection of a controlled volume of odorized air within a constant stream of desiccated air. Teflon tubing was used at and downstream from the mixing point to prevent odor lingering and cross-contamination. Odors were used at 10% vapor pressure (all PNs, 85% of KCs) or 100% vapor pressure (15% of KCs), further diluted in the dessicated air stream. We used: 1-hexen-3-ol (hx3), trans-2-hexen-1-ol (thx), cis-3-hexen-1-ol (chx), 1-hexanol (hxo), 1-heptanol (hpo), 1-octanol (oco), hexanal (hxa), heptanal (hpa), octanal (oca), nonanal (nna), 3,7-dimethyl-2,6-octadiene-nitrile (don), 3-pentanone (pnn), 2-heptanone (2hp), 3-heptanone (3hp), 5-nonanone (nnn), 6-undecanone (unn), cherry (che), mint (min), geraniol (ger), vanilla (van), citral (cit), apple (app), strawberry (str), amyl acetate (ama), benzaldehyde (bnh), methyl salicylate (mts), eugenol (eug), L-carvone (lca), D-carvone (dca), dihydro-myrcenol (dhm).

2.4.2 Tetrodes

Two types of tetrodes were used for extracellular recordings. Silicon probes were generously provided by the University of Michigan Center for Neural Communication Technology (`http://www.engin.umich.edu/facility/cnct/`). Wire tetrodes were constructed with insulated 0.0005" and 0.0004" wire (REDIOHM wire with PAC insulation). Four strands of wire were twisted together and heated to partially melt the insulation. The tip was cut with fine scissors and each channel tip was electroplated with gold solution to reduce the impedance to between 200 and 350 kΩ at 1 kHz. The same custom-built 16-channel preamplifier and amplifier were used for both types of tetrodes. Two to four tetrodes were used simultaneously. The preamp has a unitary gain, and the amplifier gain was set to 10,000×. Because of low baseline activity and low response probability in KCs (see sections 2.2.2 and 2.2.3),

fewer KCs than PNs were usually isolated in a typical recording session. Tetrodes were placed within the AL or MB soma clusters, peripheral to the neuropils at depths between 50 and 200 μm. Cell identification was unambiguous because PNs are the only spiking neurons in the locust AL, (LNs do not produce sodium action potentials, Laurent and Davidowitz, 1994), and because all the somata located above the MB calyx belong to KCs.

2.4.3 Extracellular data analysis

Tetrode recordings were analyzed as described in Pouzat et al. (2002). Briefly, data from each tetrode was acquired continuously from the four channels (15 kHz/channel, 12 bit/sample), filtered (custom-built amplifiers, band-pass 0.3–6 kHz) and stored. Events were detected on all channels as voltage peaks above a pre-set threshold (usually 2.5–3.5 times each channel's signal SD). For any detected event on any channel, the same 3 ms window (each containing 45 samples) centered on that peak was extracted from each one of the four channels in a tetrode. Each event was then represented as a 180-dimensional vector (4×45 samples). Noise properties for the recording were estimated from all the recording segments between detected events, by computing the auto- and cross- correlations of all four channels. A noise covariance matrix was computed and used for noise whitening. Events were then clustered using a modification of the expectation maximization algorithm. Because of noise whitening, clusters consisting of, and only of, all the spikes from a single source should form a Gaussian (SD = 1) distribution in 180-dimensional space. This property enabled us to perform several statistical tests to select only units that met rigorous quantitative criteria of isolation (figure 2.9).

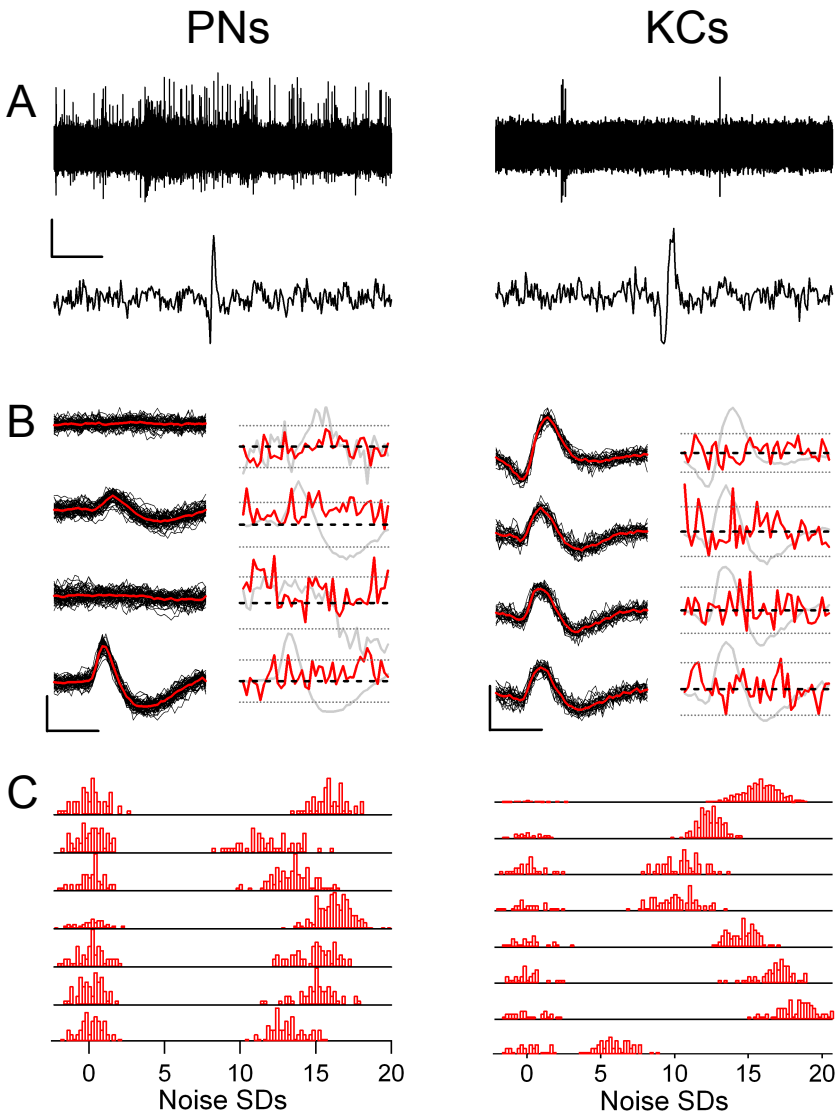

Figure 2.9. Extracellular tetrode recordings and spike-sorting. *A*, raw data traces with PN action potentials recorded in the AL (*left*), and KC action potentials recorded in the soma layer of the MB (right). Calibrations: 50 μV, 3 s (top traces), 3 ms (bottom traces). *B*, examples of two clusters: PN (left panel) and KC (right panel). In each panel the traces on the left show the superimposed events classified for that cluster (black) for each of the four tetrode channels, together with the average waveform (red). Calibration: 100 μV, 1 ms. Two of the statistical tests used to evaluate the isolation of the cells in the model are shown in *B* and in *C*: on the right side of each panel in *B* is the variance around the mean for each of the four channels, together with 95% confidence intervals which are based on the noise model. *C*, projection tests in which each pair of clusters in the model in 180-dimensional space is projected onto the line connecting the cluster centers so as to evaluate their degree of isolation. All cluster centers are separated by at least five times the noise SD. All analyzed data were selected on these separation criteria. [PN data: O. Mazor and S. Cassenaer. KC data: J. Perez-Orive.]

2.4.4 Responses

Defining what constitutes a response quantitatively and equally accurately for PNs and KCs requires careful consideration. For example, a conventional mean firing rate measure applied to the entire response period is not appropriate, because PN responses are patterned; a typical PN response, such as one composed of subsequent excitatory and inhibitory epochs, often produces a mean rate no different from baseline, and yet clearly constitutes an odor-specific response; reliability across trials thus needs to be taken into account. In addition, PNs and KCs have very different baseline firing statistics, implying that response criteria based on a change from baseline might not apply equally well to both populations. We thus analyzed the data using a variety of methods and display, in our paper, the results of one (Method A), applied identically to KCs and PNs. The analyses using other methods, summarized in table 2.1 and figure 2.10, yielded nearly identical results. Our methods are as follows. First, for all methods, we defined two response windows: short (0–1.4 s) and long (0–3 s after stimulus onset), with stimulus on for 1 s in all cases. Method A used a 3 s window. Second (Method A), a PN or KC was classified as responding during either window if its firing behavior during the window met two independent criteria of response amplitude and reliability:

1. Amplitude: The neuron's firing rate (measured in successive 200 ms bins, averaged across all trials) had to exceed n SDs of the mean baseline rate in at least one bin within the response window. Baseline rate was measured for each cell-odor pair over a period of 3 to 5 s preceding stimulus onset and over all trials with that odor. We explored values of n from 2 to 4. If n was low (e.g., $n = 2$ SDs) the rate of false responses detected in PNs prior to stimulation was

unacceptably high (>35%). If n was high (n = 4 SDs), the proportion of missed responses (as judged by visual inspection of PN rasters and PSTHs) during odor presentation was unacceptably high (>10%). Values of n of 3 or 3.5 gave low rates of both false positives (during baseline) and false negatives (during stimulation) in PNs. Values of n between (and including) 2 and 4 made no significant difference with KCs. We show the results with n = 3.5 (Method A, figure 2.4); those obtained with other values of n are summarized in table 2.1.

2. Reliability: To ensure that responses detected were reliable even at low firing rates, we required that more than half of all trials with each odor contain at least one spike during the response window. We also analyzed the same data sets using different criteria for PNs and KCs, each adapted to each population's baseline firing statistics. Despite this difference, the results (table 2.1, figure 2.7) are nearly identical to those shown in figure 2.4.

2.4.5 Sparseness

Data were analyzed using Matlab and Igor. The sparseness measures are taken from Wilmore and Tolhurst (2001), Rolls and Tovee (1995), and Vinje and Gallant (2000). In brief,

$$S_P = \frac{1 - \left(\frac{\left(\sum_{j=1}^{N} r_j / N \right)^2}{\sum_{j=1}^{N} r_j^2 / N} \right)}{1 - 1/N},$$

where N is the number of units and r_j is the response of unit j. Lifetime sparseness, S_L, is calculated in the same way, except that index j now corresponds to each odor and N to the total number of odors tested with each cell. Analog response intensities

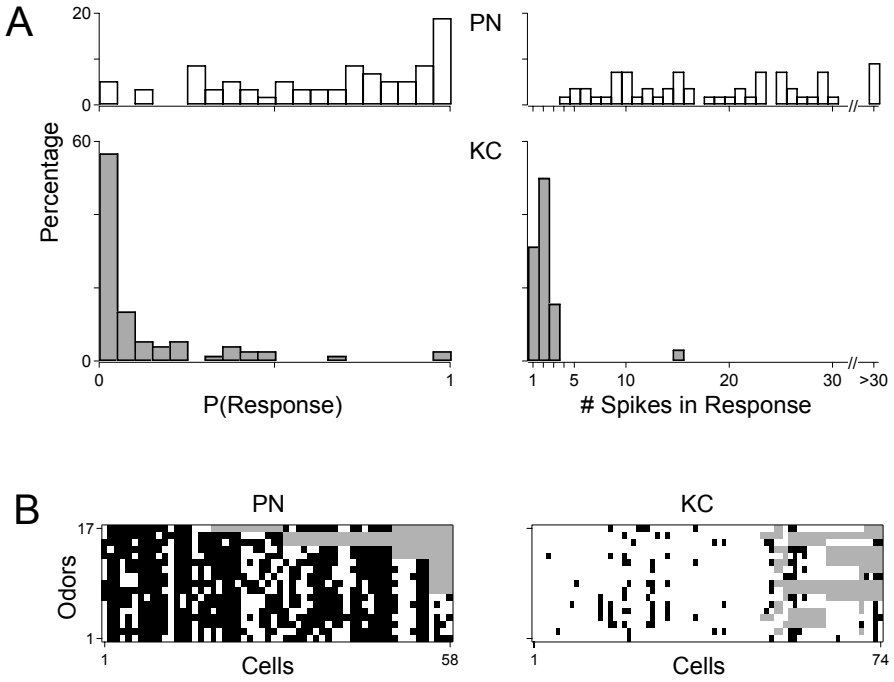

Figure 2.10. Population responses and sparseness across PNs and KCs, calculated using different criteria with PNs and KCs, for determining whether a neuron responded: A PN was qualified as responding during the 3 s following odor onset, if its firing rate increased to above 3.5 SDs of the pre-odor baseline rate (measured by a PSTH with 200 ms non-overlapping bins). In contrast, a KC response occurred when over 50% of individual trials for a particular odor showed an increase from baseline activity anywhere in the 3 s window. An increase in activity was defined as at least one 200 ms bin with a spike count higher than 3 SDs above baseline (computed from the pre-odor period over all trials). *A left*, histograms displaying PN and KC response probability distributions. Response probabilities measured across all odors tested. Note opposite skews in KC and PN distributions. *A right*, histograms displaying distributions of spike numbers in a response. Spike counts were computed only from cell-odor pairs with a significant excitatory response during the analysis window. *B*, excitatory responses (filled boxes) of individual PNs and KCs. Open squares denote inhibitory response (PNs only) or absence of a response (see figure 2.4B legend for odors). [PN data: O. Mazor and S. Cassenaer. KC data: J. Perez-Orive.]

	Method	**A**	B	C	D	E	F	G	H
	Threshold	**3.5 SD**	2 SD	3 SD	4 SD	3.5 SD			
	Window	**3 s**	3 s	3 s	3 s	1.4 s	3 s	3 s	3 s
PN	P(Resp.)	**0.64**	0.73	0.68	0.61	0.51	0.64	0.65	-
	False Pos. (%)	**2.23**	35.98	6.16	0.80	0.89	2.14	3.57	-
	Overlap (%)	**-**	90.71	96.34	96.79	87.41	99.73	99.55	-
KC	P(Resp.)	**0.11**	0.12	0.12	0.11	0.09	0.11	-	0.11
	False Pos. (%)	**0.09**	1.46	0.18	0.00	0.00	0.09	-	0.00
	Overlap (%)	**-**	99.27	99.82	99.91	97.46	100.00	-	99.46

Table 2.1. Quantitative comparison of different methods of response detection. For each method, three statistics are computed for PNs and KCs. Response probability (P(Resp.)) indicates the probability of a detected response, computed over all cell-odor pairs. The false positives value (False Pos.) is the percentage of responses detected when the method was applied to a window of baseline activity prior to odor onset (computed for all cell-odor pairs). The final statistic (Overlap), is a measure of similarity between a particular method and Method A (see section 2.4.4, p. 48), defined as the percentage of cell-odor pairs for which the two methods either both detected or both did not detect a response. Methods B–D are identical to Method A, but use a different response amplitude threshold, ranging from 2 SDs to 4 SDs above baseline. Method E is the same as Method A, but uses only a 1.4 s response window (0–1.4 s after odor onset). Method F is based on Method A, but it uses a different reliability criterion that adapts to the cell's baseline statistics. In this method, an odor response was reliable if more than half of all trials contained at least one 200 ms bin with a spike count higher than a threshold, specified as 1 SD above the mean baseline rate. Methods G and H are the methods of response detection for PNs and KCs, respectively, described in figure 2.10. [PN data: O. Mazor and S. Cassenaer. KC data: J. Perez-Orive.]

for a given cell-odor pair were computed by first segmenting the recording into 200 ms bins and computing the mean spike count in each bin, averaged over all trials with that odor. We then subtracted from all bin measures within the analysis window (1.4 or 3 s), the mean baseline rate. All so-calculated values greater than 0 over the window (7 or 15 bins) were then added. S_P and S_L vary between 0 and 1 (1 = most sparse).

2.4.6 Sharp electrode recordings and staining

Sharp electrode recordings of KCs (figures 2.5 and 2.7A) were made with borosilicate glass micropipettes (R_{DC} > 300 MΩ) filled with 0.2 or 0.5M K^+-acetate or

patch-electrode solution (see section 2.4.8). KC input resistance at the soma was usually around 1 GΩ. Intracellular recordings of LHIs (soma or dendritic impalement, figure 2.6) were made with borosilicate glass micropipettes filled with 0.5M K^+-acetate (R_{DC}:100–300 MΩ) or wire tetrodes. Intracellular staining of LHIs was carried out by iontophoretic injection of 2% neurobiotin in 0.5M K^+-acetate (0.5 s current pulses of −2.5 to −3.5 nA at 1 Hz for 30–60 min). Injected neurons were visualized in whole mounts using a diaminobenzidine-based chromogenic reaction (Wicklein and Strausfeld, 2000). Local field potentials were always recorded in the mushroom body calyx, using saline-filled patch pipettes (R_{DC}: 2–15 MΩ) or wire tetrodes. Electrical stimulation of PNs was carried out in the AL using 25 μm bipolar tungsten wires and a WPI stimulus isolator.

2.4.7 Immunocytochemistry

Partially desheathed locust brains were fixed for one hour in 5% formaldehyde, desheathed and washed for 20 h in PBS. Brains were then dehydrated through an ethanol series, placed in propylene oxide for 20 min, rehydrated and then agitated for five hours in PBS containing 5% triton and 0.5% bovine serum albumin (PBS 5% T 0.5% BSA). They were then washed for 30 min in PBS 0.5% T 0.5% BSA, and transferred to fresh PBS 0.5% T 0.5% BSA containing anti-GABA at 1:100 dilution, or, for negative control, to PBS 0.5% T 0.5% BSA lacking primary antibody. After incubation at 4°C for six days, brains were washed for two hours in PBS at room temperature and transferred to PBS 0.5% T 0.5% BSA containing fluorescein isothiocyanate-conjugated goat anti-rabbit IgG at 1:20 dilution and incubated at 4°C for four days. They were then washed for 30 min in PBS, dehydrated through ethanol series, cleared in methyl salicylate and examined by confocal laser scanning

microscopy. Figure 2.6A is a projection along the z-axis of a stack of 30 optical slices each 2.7 μm thick, constructed using the public domain *ImageJ* program (`http://rsb.info.nih.gov/ij/`). Negative control brains showed diffuse background staining.

2.4.8 Patch-clamp recordings

Whole-cell patch-clamp recordings from KCs were obtained in a semi-reduced preparation. After the brain was exposed, it was removed from the head with antenna and eyes still attached, placed on a glass coverslip in a custom-built chamber, and immobilized using insect pins placed in the eyes. The brain was then desheathed. Recordings were obtained from KC somata under visual control using a microscope with IR-DIC imaging. Patch pipettes (5–6 MΩ) were filled with a solution of (in mM): K gluconate 185, HEPES 10, EGTA 1, MgATP 4, Na3GTP 0.5 (335 mOsm, pH 7.2). Glucose (10 mM) was substituted for an equimolar amount of sucrose in the external saline solution, and the saline was bubbled continuously with O_2. Hyperpolarizing current injections (10 pA) were used to continually measure intrinsic membrane properties, and the cell was accepted for recording as long as $R_{input} > 1$ GΩ and $R_{access} < 40$ MΩ. Data was acquired on an Axopatch 1D amplifier at 10 kHz and filtered at 5 kHz. Note: In whole-cell current-clamp mode, typical EPSP duration in controls at −60 mV (figure 2.7B) was about twice that observed with sharp electrodes (figure 2.7A).

2.4.9 Picrotoxin injections

Patch pipettes were back-filled with a solution containing 1.67 mM picrotoxin and 0.3% Fast Green. After the pipette was introduced into the MB calyx (dendritic

region of the MB), a pneumatic pico-pump (WPI) was used to apply a series of four to nine 100 ms, 10 psi pressure pulses. Each pulse injected ~1 pL of solution (as measured by previous injection into a drop of oil). Injected solution remained exclusively localized to calyx, as verified by dispersal of Fast Green.

2.5 Acknowledgments

The work presented in this chapter is the product of a strong collaboration between Ofer Mazor, Javier Perez-Orive, Stijn Cassenaer, Glenn C. Turner, Rachel I. Wilson, and Gilles Laurent. O. Mazor and S. Cassenaer performed the extracellular PN recordings (figures 2.2, 2.3, 2.4, 2.9, and 2.10). O. Mazor, J. Perez-Orive and S. Cassenaer performed the quantitative comparison between PN and KC odor responses (figures 2.4, 2.10, and table 2.1). Sparseness calculations (figure 2.4) were computed by O. Mazor. Sharp electrode intracellular KC recordings were performed by J. Perez-Orive (figures 2.2, 2.3, 2.4, 2.6, 2.7, 2.8, 2.9, 2.10) and Gilles Laurent (figure 2.5). R. Wilson conducted KC whole-cell patch clamp recordings (figure 2.7). G. Turner identified and characterized the physiological and anatomical properties of the LHIs (figure 2.6), and computed the KC-LFP cross correlogram (figure 2.5).

This chapter is reprinted from *Science*, vol. 297, J. Perez-Orive, O. Mazor, G. C. Turner, S. Cassenaer, R. I. Wilson, and G. Laurent, "Oscillations and sparsening of odor representations in the mushroom body", pp 359–65, 2002, with permission.

Chapter 3

Hebbian STDP in Mushroom Bodies Facilitates the Synchronous Flow of Olfactory Information in Locusts

Odour representations in insects undergo progressive transformations and decorrelation (Mazor and Laurent, 2005; Perez-Orive et al., 2002; Wilson et al., 2004) from the receptor array to the presumed site of odour learning, the mushroom body (de Belle and Heisenberg, 1994; Dubnau et al., 2001; Yu et al., 2005; Zars et al., 2000). There, odours are represented by sparse assemblies of Kenyon cells in a large population (Perez-Orive et al., 2002). Using intracellular recordings *in vivo*, we examined transmission and plasticity at the synapse made by Kenyon cells onto downstream targets in locusts. We find that these individual synapses are excitatory and undergo hebbian spike-timing dependent plasticity (STDP) (Bi and Poo, 1998; Markram et al., 1997; Roberts and Bell, 2002) on a ± 25 ms timescale. When placed in the context of odour-evoked Kenyon cell activity (a 20- Hz oscillatory population discharge), this form of STDP enhances the synchronization of the Kenyon cells' targets and thus helps preserve the propagation of the odour-specific codes through the olfactory system.

3.1 Introduction

Olfactory processing in insects begins in an array of receptor neurons that express collectively many tens of olfactory receptor genes (~60 in Drosophila (Clyne et al., 1999; Vosshall et al., 2000); ~150 in honeybees (The Honeybee Genome Consortium, 2006). The representations of general odours are then decorrelated by local circuits of projection neurons and local neurons in the antennal lobe (Mazor and Laurent, 2005; Perez-Orive et al., 2002; Wilson et al., 2004). In locusts and other insects, the antennal lobe output is distributed in space and time and can be described as stimulus-specific time-series of projection- neuron activity vectors, updated at each cycle of a 20-Hz collective oscillation (Mazor and Laurent, 2005; MacLeod and Laurent, 1996; Wehr and Laurent, 1996). Distributed projection-neuron activity is then projected to Kenyon cells, the intrinsic neurons of the mushroom body. In contrast to projection neurons, Kenyon cells respond very specifically and fire extremely rarely (Perez-Orive et al., 2002). The mechanisms underlying this sparsening are starting to be understood (Perez-Orive et al., 2002; Jortner et al., 2007). Such sparse representations are advantageous for memory and recall (Jortner et al., 2007), consistent with established roles of the mushroom bodies in learning (de Belle and Heisenberg, 1994; Dubnau et al., 2001; Yu et al., 2005; Zars et al., 2000). In Drosophila, experiments combining molecular inactivation with behaviour indicate that synaptic output from Kenyon cells in the lobes is required for memory retrieval (Dubnau et al., 2001). Little is known, however, about the electrophysiological properties of these synapses.

3.2 Results

3.2.1 Synaptic connections between individual Kenyon cells and β-LNs

We studied the connections made by Kenyon cells onto a small population of extrinsic neurons (MacLeod et al., 1998) in the β-lobe of the locust mushroom body (figure 3.1a), using an intact, *in vivo* preparation (Methods).

β-lobe neurons (β-LNs) respond to odours; their responses are odour-specific and their tuning is sensitive to input synchrony (MacLeod et al., 1998). We recorded intracellularly from pairs of Kenyon cells and β-LNs: randomly selected Kenyon cells were impaled in their soma; β-LNs were impaled in a dendrite in the β-lobe. We focused on one β-LN anatomical subtype (MacLeod et al., 1998), which comprises many individual neurons. Neurons of this subtype, called β-LNs here, could be recognized also by their physiological characteristics (see below). Each β-LN has extensive dendrites (figure 3.1a, and figure 3.5) that intersect many of 50,000 Kenyon cell axons. Monosynaptic connections were found in ~2% of tested Kenyon cell KC–β-LN pairs (figure 3.1b). All were excitatory. The delay between Kenyon cell spike and β-LN excitatory post-synaptic potential (EPSP) onset was 6.5 ± 0.70 ms, including 5.4 ± 0.25 ms for spike propagation from Kenyon cell soma to the β-lobe. The remaining (synaptic) delay (~1 ms) is similar to that at another chemical synapse in the locust brain (Jortner et al., 2007). Unitary EPSPs were large (1.58mV ± 1.11, n = 9 pairs), in contrast to those generated in Kenyon cells by individual projection neurons (86 μV ± 44) (Jortner et al., 2007). The fact that Kenyon cell outputs are powerful is consistent with Kenyon cell spikes being rare and therefore highly informative. EPSP amplitude varied greatly across connected pairs (0.55−4 mV). This could reflect a distribution of electrotonic distances between synapses

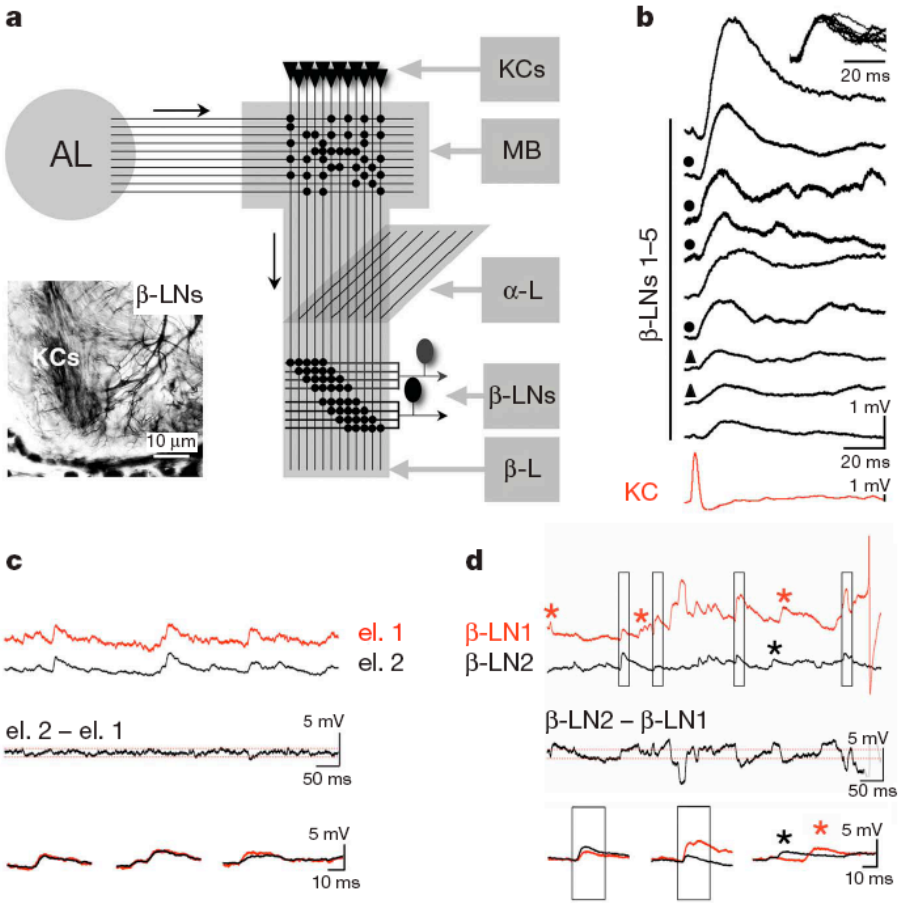

Figure 3.1. Synaptic connections between individual Kenyon cells and β-LNs are excitatory, powerful and varied in gain. a, Schematic of the locust olfactory circuits: Projection neuron axons (n = 830) exit the antennal lobe (AL) and send collaterals into the mushroom body (MB) calyx. There, they excite Kenyon cells (n = 50,000) with, 50% average connectivity (Jortner et al., 2007). Kenyon cells each send a bifurcating axon into the α- and β-lobes (α-, β-L), forming 'beams' of thousands of tightly packed axons (inset). The finer dendrites of β-LNs run normal to Kenyon cell axons in β-LN-specific sectors: two of these can be seen in the inset (a photomicrograph of a 10-μm-thick section of the distal end of the β-lobe). **b**, Spike-triggered averages of β-LN intra-dendritic recordings from nine different KC–β-LN pairs (9 Kenyon cells, 5 β-LNs). All β-LNs and Kenyon cells are recorded *in vivo*, with intracellular electrodes. Note the wide-range of spike-triggered average amplitudes. Same-symbol marked spike-triggered averages are from the same β-LNs, with different presynaptic Kenyon cells sampled successively. Inset, scaled spike-triggered averages in **b**, illustrating similarity of kinetics. **c**, Simultaneous dendritic impalements of one β-LN with two separate electrodes (el. 1, 2). el.2−el.1 is the difference between the two voltage traces; note the high correlation of amplitudes (noise envelope is 2 s.d., red stipples), indicating similar electrotonic access to synaptic sites. Lower panel, overlay of selected EPSPs from above. **d**, Simultaneous dendritic impalements of two different β-LNs (1 and 2). Note some common EPSPs (boxes) and EPSPs specific to either β-LN (*). β-LN2−β-LN1 is the difference between the two voltage traces; note significant variations on each side of the noise envelope (stippled lines as in panel **c**). Lower panel, overlay of selected EPSPs from above.

and recording sites. Simultaneous impalements of different dendrites in the same β-LN (n = 2 experiments), however, show that the amplitudes of most events were the same across recording sites (PearsonÕs correlation >0.9) (figure 3.1c). Consistent with this, unitary EPSP kinetics (10-90% rise time, 8.3 ms ± 2.3; time to 1-(1/e) of peak, 13.2 ms ± 4.4) were independent of the β-LN recorded and, thus, of the impalement site (inset, figure 3.1b). Simultaneous dendritic recordings of different β-LNs (n = 5 experiments), however, revealed that their synaptic backgrounds overlapped only partly (figure 3.1d; Pearson's correlation or fraction of common EPSPs, 0.1Ð0.3). Common EPSPs rarely had the same amplitude (figure 3.1d). Hence, β-LNs may each receive inputs from hundreds to thousands (~2% of 50,000 Kenyon cells) of Kenyon cells, in overlapping subsets; KC–β-LN connections are strong on average, with target-specific strength.

3.2.2 β-LN tuning and spike-time precision during responses to odours

Odour-evoked activity in projection neurons and Kenyon cells consists principally of sequential volleys of synchronized spikes– generally, one spike per responding neuron per oscillation cycle (Mazor and Laurent, 2005; Perez-Orive et al., 2002; Wehr and Laurent, 1996). β-LN responses to odours also consisted typically of sequences of single phase-locked spikes, timed around the trough of several local field potential (LFP) oscillation cycles (figure 3.2a, b). The cycles when a spike was produced (usually with probability <1) depended on β-LN and stimulus identity, as illustrated in (figure 3.2c, d. We conclude that, to each oscillation cycle corresponds a particular activity vector in the projection neuron (Perez-Orive et al., 2002), Kenyon cell (Perez-Orive et al., 2002) and β-LN populations. By recording from pairs of β-LNs simultaneously during odour trials, we also observed that, when the two β-LNs

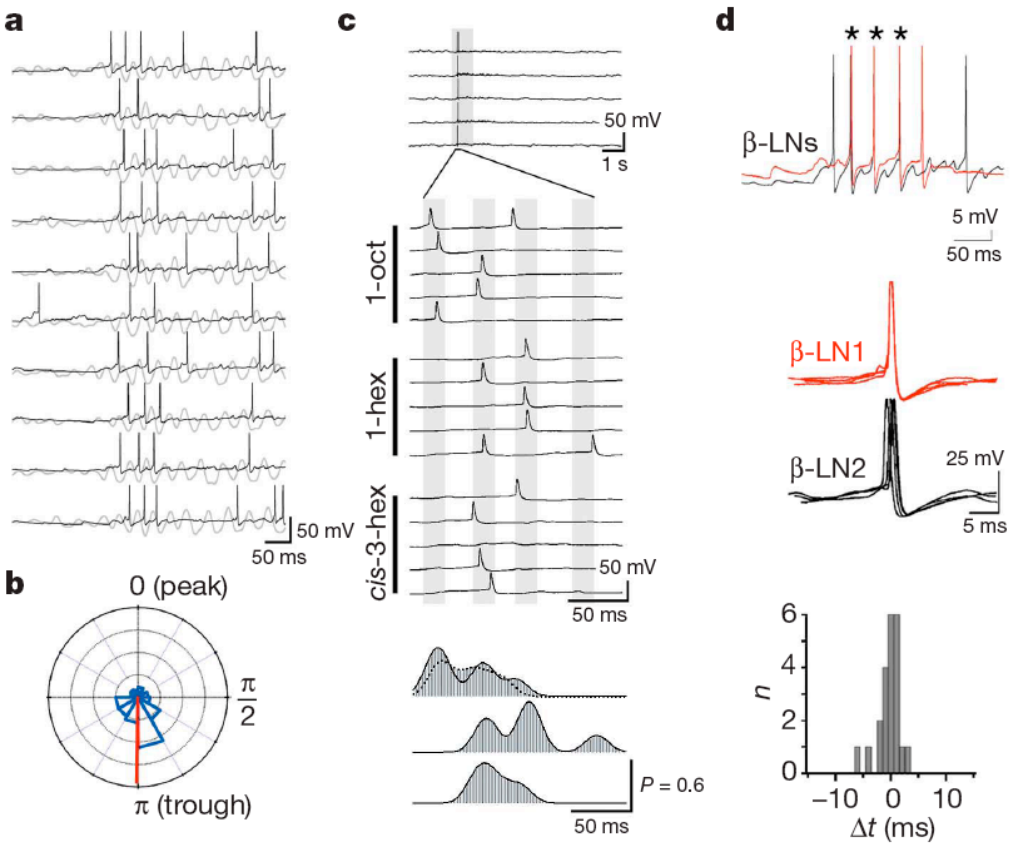

Figure 3.2. *β*-**LN tuning and spike-time precision during responses to odours. a**, Responses of one *β*-LN to ten successive trials with odour cis-3- hexen-1-ol (LFP shown in grey). Note non-random timing of *β*-LN action potentials during each trial and during each oscillation cycle (calibration LFP, 250 μV).**b**, *β*-LN action potentials lock to the trough of LFP during odour responses. Phase plot of *β*-LN action potentials during odour responses (phases plotted clockwise). Distribution (blue) constructed with data from 8 *β*-LNs (average in red). **c**, *β*-LN responses are structured in time and are odour-specific (see also (MacLeod et al., 1998)). This *β*-LN responds to different odours with different discharge patterns: upper panel, response to odour 1-octanol (grey bar); middle panels, same 5 trials as in upper panel on an expanded time base, showing approximate timing of relevant oscillation cycles (grey bars) and also showing responses to 1-hexanol and cis-3-hexen- 1-ol. Apparent spike-time jitter is due to variability of oscillation cycle duration within and across trials. In contrast, spike-time jitter across *β*-LNs, but within the same oscillation cycle of the same trial, is low (see panel **d**). Lower panel, smoothed peri-stimulus time-histogram from the trials and odours in the middle panels (in the same order). The y axis measures firing probability. Stippled line in the top peri-stimulus time-histogram shows responses to a second set of 5 trials with this odour, delivered after trials with the 2nd and 3rd odours.**d** Spike discharges in response to odours are precisely locked across *β*-LNs. Upper panel, simultaneous recording of two *β*-LNs during odour stimulus, with spikes from both neurons in several oscillation cycles (*). Note precise overlap. Middle panels, zoom on those action potentials, superimposed and triggered on *β*-LN1 action potential. Lower panel, distribution of spike-time jitter. Recordings as in the middle panels; Δt is the time difference between spikes in *β*-LNs 1 and 2, when they occur in the same oscillation cycle.

fired one action potential during the same oscillation cycle (n = 4 pairs; (figure 3.2d, upper panel, asterisks), those action potentials were tightly synchronized (±2 ms, (figure 3.2d, middle and lower panels).

3.2.3 Hebbian spike-timing-dependent plasticity at the KC-β-LN synapse

A fortuitous observation provided hints of plasticity at the KCÐβ- LN synapse (figure 3.3a). At trial 4 of a Kenyon cell stimulus sequence intended to explore β-LN integration, the β-LN fired a spontaneous action potential roughly at the time of the first (of 2) Kenyon-cell evoked EPSP (figure 3.3a). At trial 5, 10 seconds after this single fortuitous pairing, the first EPSP of the pair was greatly enhanced (figure 3.3a). This suggested the possibility of spike-timing-dependent plasticity (STDP), a phenomenon thus far unknown in invertebrates but well characterized in vertebrates, in which the gain of a connection can be changed according to the temporal relationship between pre- and post-synaptic spikes (Bi and Poo, 1998; Markram et al., 1997; Roberts and Bell, 2002). We explored the consequence of preÐpost temporal relationships on the KC–β-LN synapse. A β-LN was impaled and stimulated alternately by two independent Kenyon cell pathways–one for pairing, one for unpaired control (figure 3.3b). Each stimulus was repeated every 10 s, with a 5-s delay between pairing and control stimuli. Pairing consisted of a single Kenyon cell (pre) stimulus and a 5-ms supra-threshold β-LN (post) current pulse, timed such that the delay ($dt = t_{post} - t_{pre}$) between pre- and post-synaptic spikes varied between −60 and +50 ms. Test trials, used to measure connection strength before and after pairing, were identical to the pairing trials in all respects except in the temporal relationship between pre- and post-synaptic spike times (2.5 s apart, (figure 3.3b). Two examples with controls are shown in Fig. 3c, d (for dt = 10 ms and

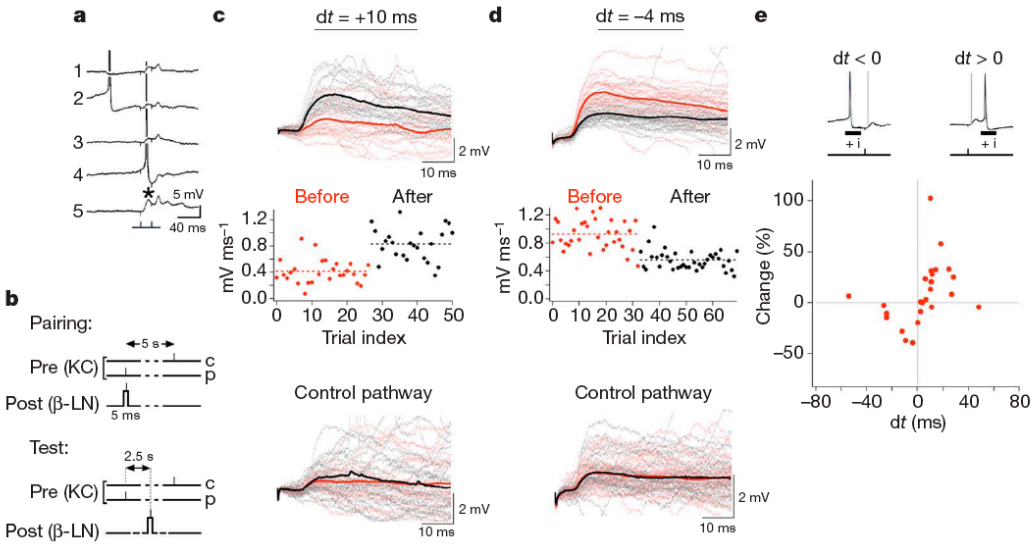

Figure 3.3. Hebbian spike-time-dependent plasticity at the KC–β-LN synapse. a, Effect of single, near-coincident spikes in pre- and post-synaptic neurons on the Kenyon cell to β-LN synapse. The β-LN is held from a dendrite at normal resting potential *in vivo*, and subjected to paired Kenyon cell stimuli at 20-ms intervals (first and second stimuli in train are from different Kenyon cell electrodes). At trial 4, the effect of the first stimulus sums with an on-going slow depolarization, causing a β-LN action potential. At trial 5, the first EPSP (asterisk) is dramatically enhanced. **b**, Stimulation protocols to probe STDP. For 'pairing', the β-LN is depolarized with a 5-ms DC pulse, causing a single β-LN spike in a window around the time of the Kenyon cell pairing stimulus (p). We tested the effects of 5 to 25 successive pairings (10-s intervals); all evoked STDP. Control Kenyon cell stimulus (c) is offset by 5 s relative to the β-LN current pulse. 'Test', same protocol as pairing, except that Kenyon cell and β-LN stimuli are 2.5 s apart. **c**, Potentiation of KC–β-LN connection after 25 pairings with dt = 10 ms. Upper panel, superimposed before- and after-trials, with their averages (bold lines). Middle panel, EPSP slope against trial index, with 25 pairings between the before- and after- periods. Stippled lines are average slopes over corresponding trials. Lower panel, a control, recorded at the same time with a second Kenyon cell pathway, offset by 5 s with respect to the β-LN spike, and showing no significant change. **d**, Depression of KC–β-LN connection after 25 pairings with dt = −4 ms. Panels as in panel **c**.**e**, STDP plot for 26 values of dt = t_{post}−t_{pre}, where dt is measured as the delay between the β-LN spike (t_{post} caused by intracellular current injection, +i) and the β-LN EPSP onset (t_{pre}, grey line, upper panels). EPSP onset time is used (rather than Kenyon cell stimulus time) because Kenyon cell spike time at the β-LN synapse is delayed from stimulus time, owing to spike propagation (see (figure 3.4a). Note the zero-crossing is slightly offset from dt = 0 (lower panel).

−4 ms, 25 pairings each). For dt = 10 ms (figure 3.3c), the paired input underwent potentiation; for dt = − 4 ms (figure 3.3d), it underwent depression. For both conditions, the control pathway (same β-LN, different Kenyon cell input) remained unchanged (figure 3.3c, d, lower panels). The changes were thus input-specific; they were often detectable after a single pairing (see also (figure 3.3a), and could be maintained for up to 25 min. We tested 26 values of dt between −60 and +50 ms. The resulting changes (figure 3.3e) define a classical hebbian profile (Bi and Poo, 1998; Roberts and Bell, 2002): the synapse is potentiated when pre- precedes post-, and depressed when post- precedes pre-, with symmetrical profiles. The changes could be fitted well with two exponential decays flanking a narrow linear range around t = +4ms (τ_1 =10.4 ms for dt < −9 ms; y = 3.78t −13.1 for −9ms < dt <17.5 ms; τ_2 =11.6 ms for dt >17.5 ms). Several connections were tested successively with two (or more) values of dt (some positive, others negative): the same connections could undergo both depression and potentiation, depending on the value of dt. The STDP profile thus seems to be a property of each connection and not only a collective one.

3.2.4 The effect of STDP on β-LN spike timing

We observed that the values of dt over which synaptic weights change correspond to the period of single odour-evoked oscillation cycles; hence, only within-cycle 'coincidences' may modify the connections between a Kenyon cell and its targets. The features of the STDP curve, when considered together with the timing of Kenyon cells and β-LNs during odour-evoked activity, have interesting consequences. Consider the phases of Kenyon cell and β-LN spikes (figure 3.4a). Owing to propagation delays, Kenyon cell spikes reach their targets just before the trough of the LFP, a

little before β-LN firing (figure 3.2a, b). Consider a cycle in which a β-LN spikes early (dt <0): some KC–β-LN connections will undergo depression (figure 3.3e); at the next trial, β-LN spike time at this cycle should be delayed (figure 3.4b). If, in contrast, a β-LN spikes late, STDP should potentiate Kenyon cell drive to it, and thus advance spike time for that cycle (figure 3.4b). In short, the cycle-by-cycle action of STDP suggests adaptive control of β-LN spike phase. The need for such regulation is not unique to this system: models of cortical networks indicate that, as activity propagates through successive 'layers', accumulating noise can rapidly smear the temporal structure that may exist (Diesmann et al., 1999; Vogels et al., 2005). Modelling studies (Arthur and Boahen, 2006; Suri and Sejnowski, 2002; Zhigulin et al., 2003) predict that STDP, given appropriate parameters, could preserve the temporal discretization of activity through such layers.

We generated a reduced model of the KC–β-LN circuit (Methods) and introduced the STDP rule derived from our experiments (figure 3.4c). To control the relative phases of Kenyon cells and β-LNs, we drew Kenyon cell spike phases from experiments (Perez-Orive et al., 2002), and input weights from uniform distributions with different means: with low weights, β-LN spikes tended to occur late (dt >0, (figure 3.4d); with larger weights, they occurred early (dt < 0, figure 3.4d). After several trials (each with a random draw of inputs from the same distribution), STDP was allowed to modify synaptic weights for the following trials: when β-LN spikes occurred late (dt > 0), Kenyon cell outputs became potentiated and β-LN spikes were advanced; for dt < 0, time shifts were inverted. The histograms in figure 3.4d represent spike-time distributions for 1,000 trials before (red) and after (black) STDP, for each of three conditions. These simulations were repeated 200 times (50 trials each), with 11 different Kenyon cell input distributions (figure 3.4e).

Once STDP was turned on (trial 1), the evolution was systematic and rapid, leading to the adaptive up- or downregulation of input weights, firing phase and response intensity (top, middle and bottom, respectively, all averages; figure 3.4e). Given that the model is entirely constrained by experiments, it is noteworthy that the mean phase of the first β-LN spike at steady state (π rad, figure 3.4e), matches precisely that measured experimentally (figure 3.2b).

Figure 3.4 ***(on the next page).*** **The effect of STDP on β-LN spike timing. a**, Polar plot of Kenyon cell spike phase in the calyx (somata) (KC_C) and in the β-lobe (KC_L), and β-LN spike phase (from dendrites in the β-lobe) relative to the LFP (in the calyx). All measurements come from experiments. Green and red lines indicate extrema of the STDP curve (see panel **c**). **b**, Schematic of temporal relationships between LFP, Kenyon cell spike time, β-LN spike time and the STDP rule. The Kenyon cell mean spike time in the calyx (KC_C) is about $\pi/2$ after the LFP peak, and near the LFP trough in the β-lobe (KC_L) , owing to propagation delay. β-LN mean spike time in the β-lobe is at the LFP trough (π rad). The STDP curve is represented in colour gradients. The predicted effect of STDP on β-LN spike time is schematized underneath. If the β-LN spike occurs early, STDP should depress late Kenyon cell inputs (in this oscillation cycle), delaying this β-LN spike at the next opportunity. The converse applies if the β-LN spike occurs late. **c**, The STDP fit (two exponentials flanking a linear segment, see text) overlaid on experimental data. **d**, Simulations of STDP on β-LN spike time (rasters) (model β-LN excited by 10 model Kenyon cells during one LFP cycle). First trial at top. Three conditions are illustrated: left panel, low input weights (mean, 1.8 mV; range, 3 mV), causing late β-LN-spike times (dt > 0); when STDP is turned on, potentiation shifts β-LN spikes to earlier times. Right panel, high input weights (mean, 9 mV; range, 3 mV); when STDP is turned on, depression delays β-LN-spike times. Middle panel, intermediate weights, causing no change. Histograms show the distribution of β-LN spike times before (red) and after (black) STDP (1,000 runs per condition). **e**, Evolution of KC–β-LN weights (upper panel), β-LN mean spike phase (middle panel) and number of spikes per response (lower panel) over 50 trials following onset of STDP (at trial 1). AP, action potential. Each curve is an average of 200 simulations (11 different input distribution means). Asterisks indicate a corresponding condition in the three plots. **f**, Schematic of experimental design. LFP cycles (1...k) during which the recorded β-LN-produced action potentials are selected. Peaks of LFP are used to trigger a sequence of current (I_m) pulses (a, b, c) into the β-LN (one such sequence per oscillation cycle). The a-b-c sequence lasts less than one oscillation cycle, and is repeated for all selected cycles, over several trials. **g**, Example of protocol described in panel **f**, such that the β-LN spike is phase-delayed to ~3/2 π (arrows). (Interrupted segments of β-LN potential trace are bridge-balance artefacts.) **h**, Phase plot of spikes in one β-LN before and after pairing (10 trials each), as in panel **g**. Fifteen pairing trials (estimated dt = 17 ms). **i**, Distributions of pairing induced mean phase shifts, measured in 20 separate experiments (6 β-LNs; mean ± s.e.m = − 0.74 ± 0.4 ms versus 2.0 ± 0.4 ms).

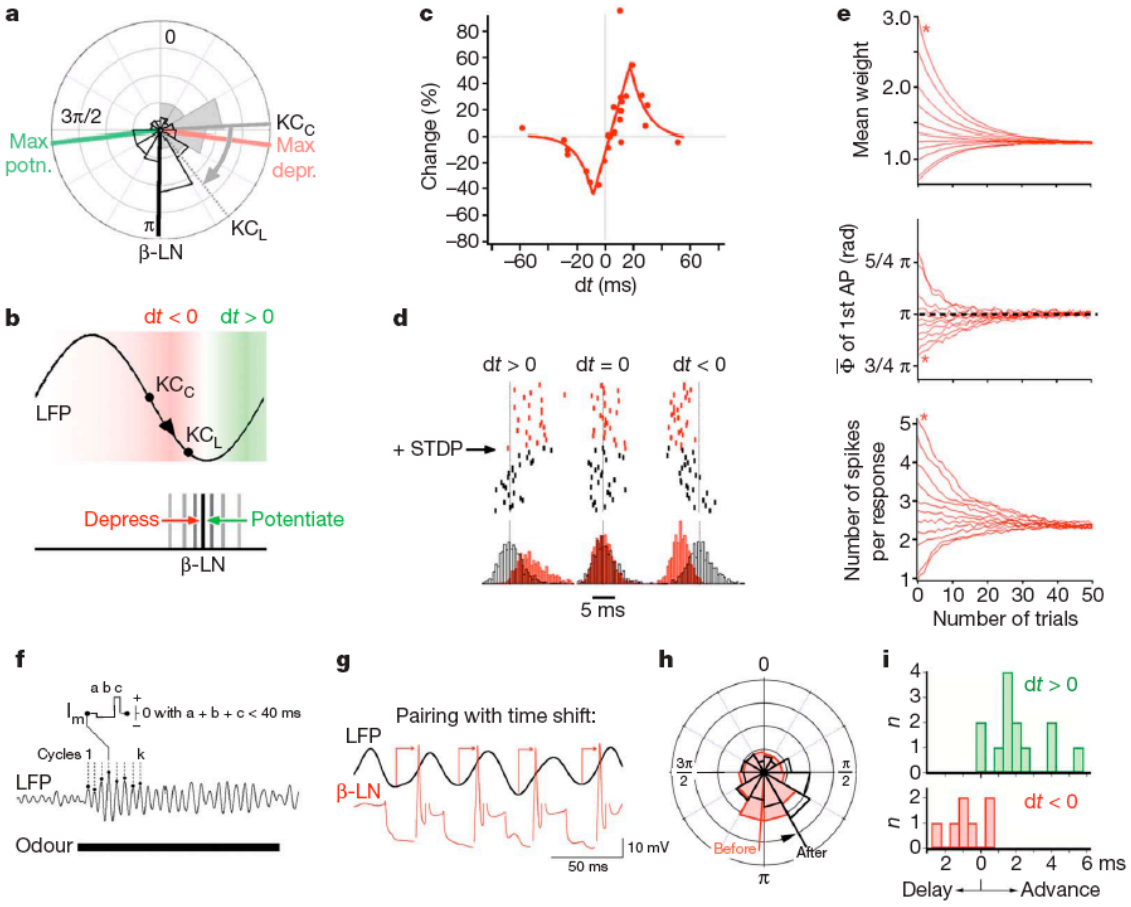
a
0
3π/2
Max potn.
KC_C
Max depr.
π
KC_L
β-LN
b
dt < 0
dt > 0
KC_C
LFP
KC_L
Depress
Potentiate
β-LN
c
Change (%)
dt (ms)
d
dt > 0
dt = 0
dt < 0
+ STDP
5 ms
e
Mean weight
Φ of 1st AP (rad)
5/4 π
π
3/4 π
Number of spikes per response
Number of trials
f
a b c
I_m
0 with a + b + c < 40 ms
Cycles 1
k
LFP
Odour
g
Pairing with time shift:
LFP
β-LN
10 mV
50 ms
h
0
π/2
π
3π/2
Before
After
i
dt > 0
dt < 0
n
Delay
Advance
6 ms

To test directly the effect of STDP on β-LN output, we next manipulated β-LN spike timing during responses to odours *in vivo*: if our model is correct, such manipulations should change the output of the odour-activated Kenyon cells onto that β-LN and, thus, generate predictable shifts in its spike phase. During odour stimuli, short current pulses locked to selected cycles of the LFP were injected in a β-LN: a negative pulse (b, figure 3.4f) was injected during the cycles and phase when the β-LN would naturally fire (to prevent stimulus evoked spikes), and a positive pulse (c, figure 3.4f) was injected at a desired phase, for those same cycles (that is, at an abnormal time relative to the Kenyon cell inputs that would normally drive the recorded β-LN). An example is shown for four consecutive cycles in figure 3.4g. After several such pairing trials, current injection was terminated and β-LN-firing phase over the next trials was compared to that before pairing. figure 3.4h plots the effects of one such manipulation ($dt > 0$): as predicted, an artificial phase-delay caused a corrective phase-advance. Twenty distinct experiments were carried out in six β-LNs; the expected phase shifts were observed in 16 of those 20 (Mann-Whitney: $P < 0.001$) (figure 3.4i). This is consistent with an adaptive role for STDP in the fine-tuning of β-LN spike-phase, and may explain the tight synchronization of β-LNs (figure 3.2d). Hence, STDP helps preserve the discrete and periodic structure of olfactory representations as they flow through the mushroom bodies.

3.3 Discussion

We showed that the connections made by Kenyon cells to β-LNs are excitatory, strong on average, variable across pairs, and plastic. Plasticity follows time-sensitive hebbian associativity rules (Bi and Poo, 1998; Markram et al., 1997; Roberts and Bell, 2002) and is constrained to within-cycle interactions between pre-and postsynaptic

neurons. STDP is therefore not specific to vertebrates or cortical architectures. We do not know the molecular underpinnings of STDP in this system, or whether STDP might confer the associative features usually ascribed to mushroom bodies (de Belle and Heisenberg, 1994; Dubnau et al., 2001; Yu et al., 2005; Zars et al., 2000). The fly and honeybee genomes both reveal coding sequences for N-methyl-D-aspartate (NMDA) receptor subunits (The Honeybee Genome Consortium, 2006; Ultsch et al., 1993) and some Drosophila behavioural results (Tanimoto et al., 2004) are compatible with STDP learning rules (Drew and Abbott, 2006). One hypothesis, readily testable here, is that STDP provides associativity by tagging transiently the subset of synapses activated simultaneously by the odour, before the conditional arrival of a slower, non-specific reward signal (Frey and Morris, 1998).

Our results reinforce the proposed importance of spike timing for this, and possibly other, olfactory system(s) (Mazor and Laurent, 2005; Perez-Orive et al., 2002): Kenyon cells act as coincidence detectors for synchronized projection neuron input (Perez-Orive et al., 2002), β-LNs act as coincidence detectors for Kenyon cell input; because STDP helps enhance β-LN synchronization, we infer that spike timing must be relevant also for the processing of β-LN output. These results indicate that the oscillation cycle – a temporal unit of processing first defined by negative feedback in the antennal lobe (MacLeod and Laurent, 1996) – is actively preserved in at least three successive layers of processing (projection neurons, Kenyon cells and β-LNs). It will be interesting to assess whether all Kenyon cell outputs obey the same STDP rules, and if these rules are themselves subject to learning related modifications. Indeed, Kenyon cells seem to communicate with one another through axo-axonal chemical synapses (Leitch and Laurent, 1996). Given the dynamics of projection neuron/Kenyon cell activity vectors in response to odours (Mazor and Laurent, 2005;

Perez-Orive et al., 2002; Wehr and Laurent, 1996) the possibility that Kenyon cell–Kenyon cell synapses also undergo STDP suggests a mechanism for sequence learning (Nowotny et al., 2003), similar to principles proposed for spatial map formation in rodents (Blum and Abbott, 1996; Mehta et al., 2002); here, however, the learned sequences have no relation to movement in physical space. The existence of such similarities (synaptic learning rules, and synchronized and sequential neural activity patterns) may bring us closer to understanding the relationships between circuit dynamics, architecture and learning in the brain.

3.4 Methods

3.4.1 Preparation and stimuli

All results were obtained *in vivo* from locusts (*Schistocerca americana*) in an established, crowded colony. Young adults of either sex were immobilized, with two antennae intact for olfactory stimulation. The brain was exposed and de-sheathed as previously described (Perez-Orive et al., 2002). Odours were delivered by injection of a controlled volume of odourized air within a constant stream of dessicated air. Teflon tubing was used at and downstream from the mixing point to prevent odour lingering and cross-contamination. Odours were used at 10% vapour pressure further diluted in the dessicated air stream. The results presented here originate from recordings of over 50 β-LNs in 40 locusts.

3.4.2 Electrical stimulation

Twisted-wire tetrodes obtained from FHC (number CE4B75) were modified for monopolar stimulation, with the casing serving as the anode. The tips of the tetrodes

were splayed such that the distance between the exposed tips was approximately equal to 60% of the diameter of the mushroom body calyx. The exposed end of the stimulating electrode was embedded among Kenyon cell somata. The tetrodes were electroplated with gold solution to reduce the impedance to between 200 and 350 kΩ at 1 kHz. Stimulating currents (5–50 μA, 0.1 ms) were generated by an STG1000 Multichannel System. The number of consecutive pairing trials varied between 5 and 25, at 10-s intervals. Propagation delays for the Kenyon cell action potentials were measured as the delay between a Kenyon cell soma stimulus and the extracellular spike volley recorded in the β-lobe, at the level of the β-LN dendritic recordings (Perez-Orive et al., 2004).

3.4.3 Intracellular recordings

Sharp electrode recordings from the dendrites of β-lobe neurons were made with borosilicate glass micropipettes (DC resistance, 100MΩ) filled with 3M K acetate. Input resistance was around 300MΩ. The cell type from which the data are derived could be recognized by several characteristics, including response to odour, sub-threshold baseline activity profile, and response to electrical stimulation of Kenyon cells. A series of pilot experiments, in which the cells were stained intracellularly by injection of 6% cobalt hexamine, confirmed that cells with these physiological characteristics belong to a specific morphological class (figure 3.5). EPSP slopes were measured from linear fits to voltage trace between 10% and 90% of rising phase. Recordings from β-LNs were always made from dendrites in the β-lobe (the largest dendrites are often several μm in diameter). That these recordings were not from Kenyon cell axons is guaranteed by the fact that Kenyon cell axons are too small for intracellular impalement (100–400nm diameter (Leitch and Laurent,

1996)). This identity of β-LNs was confirmed by dye injection. Kenyon cell intracellular recordings were always made from their somata (5–7 μm diameter).

3.4.4 Field recordings

Twisted-wire tetrodes obtained from FHC (number CE4B75) were used for extracellular recordings of the local field potential (LFP). For these recordings, the tip was cut with fine scissors and each channel tip was electroplated with gold solution to reduce the impedance to between 200 and 350 kΩ at 1 kHz. These recordings were made with a custom-built 16-channel preamplifier and amplifier. Two to four tetrodes were used simultaneously. The pre-amp has a unitary gain, and the amplifier gain was set to 10,000X. For pairing experiments during odour stimuli (figure 3.4fÐi), the LFP was low-pass filtered on line and fed through a real-time peak-detection algorithm. Each detected peak was given a rank order (1...k) and the cycles during which the recorded β-LN produced an action potential were identified. Because of intertrial variability of LFP and small uncertainty about cycle ranking in each trial, we typically selected eight consecutive oscillation cycles centred on the cycles of interest. During each one of those 8 consecutive cycles, a current-pulse sequence (a–c, figure 3.4f) was injected into the β-LN dendrite so as to phase-advance ($b > 0$; $c < 0$) or phase-delay ($b < 0$; $c > 0$) its odour-evoked spikes. This pulse sequence was started again at each LFP peak of the selected cycles.

3.4.5 Simulations

A leaky integrate-and-fire model (R = 300M Ω, τ = 20 ms) was implemented in Igor (Wavemetrics, Lake Oswego, Oregon). Spiking threshold and after-hyperpolarization were estimated from intracellular recordings. A single oscillation cycle was mod-

elled as 10 inputs convolved with a current waveform derived from dual KC–β-LN recordings. The number of inputs per cycle was based on KC–β-LN connectivity estimates from dual intracellular recordings and from β-LN baseline sub-threshold activity, as well as on the average Kenyon cell population response time course (Mazor and Laurent, 2005; Perez-Orive et al., 2002; Wilson et al., 2004) and response probability (Perez-Orive et al., 2002). The specific Kenyon cell spike times were drawn randomly at every trial from the Kenyon cell spike time distribution. The weights at the start of each simulation were drawn randomly from a distribution similar to that observed experimentally, and scaled for different simulations to evaluate the effect on the spiking response in the presence and absence of STDP. The STDP rule was modelled, as described in the text, by two exponentials flanking a linear region (fitted to data by minimizing Chi-square; Igor curve-fitting). When STDP was invoked, every trial was followed by an update of the weights, on the basis of dt, as dictated by the fitted STDP curve (figure 3.4).

3.5 Acknowledgements

This work was supported by an NIH training grant, grants from the NIDCD, and the Lawrence Hanson Fund. We thank E. Schuman, I. Fiete, M. Murthy, M. Papadopoulou, O. Mazor, V. Jayaraman and the reviewers for their helpful comments.

This chapter is reprinted from *Nature*, vol. 448(7154), S. Cassenaer, and G. Laurent, "Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts', pp 709–13, 2007, with permission.

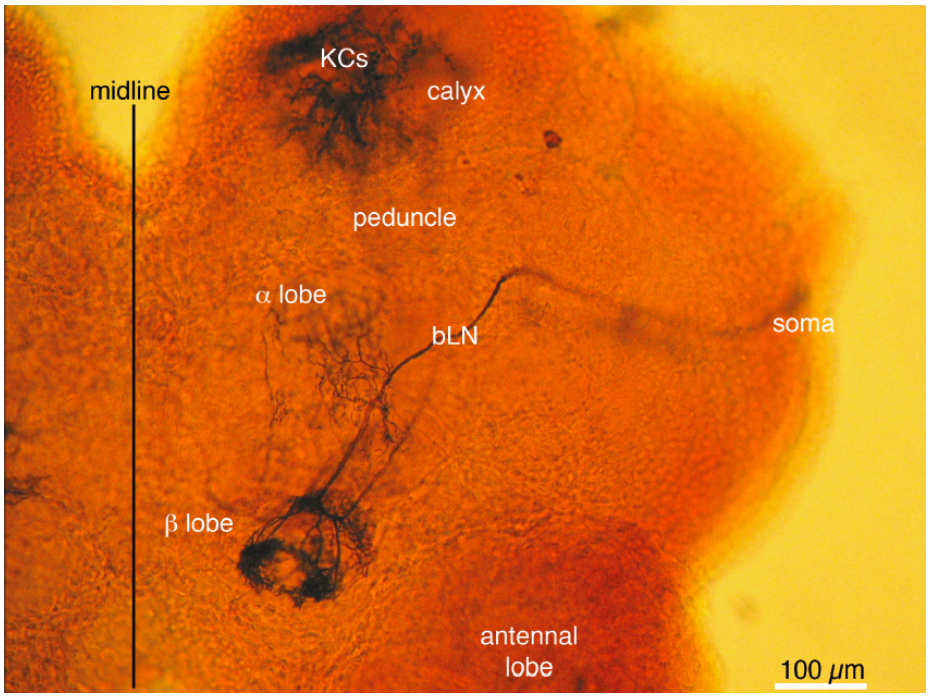

Figure 3.5. Wholemount of an intracellular fill (cobalt hexamine, TimmÕs silver intensification) of a β-LN of the type examined here. Dorsal is to the top, lateral to the right. Some Kenyon cells were also filled (after the β-LN) in this mushroom body. (The calyx, peduncle and lobes together form the mushroom body.)The β-LN was impaled in a dendrite, in the β lobe. The KCs were impaled in their soma. Only the KC dendrites are visible here. The top half of the roughly spherical antennal lobe is seen at the bottom of the figure. The β lobe runs at 45 degree along the medial edge of the antennal lobe.

Chapter 4

Concluding Remarks

4.1 Transformation of Odor Representations

Odorants bind to olfactory receptors of primary sensory neurons (olfactory receptor neurons, ORNs) in the insect antenna, causing the ORNs to fire and, in turn, activate projection neurons (PNs) in the antennal lobe (AL). This gives rise to odor- and PN-specific activity patterns, which are locked to a globally coherent 20-30 Hz oscillation shaped by local inhibitory neurons (LNs) (Wehr and Laurent, 1996; MacLeod and Laurent, 1996; Stopfer et al., 1997; MacLeod et al., 1998). The third-order olfactory neurons of the insect brain, contacted by the PNs, are the Kenyon cells (KCs, Kenyon, 1896; Laurent and Naraghi, 1994).

4.1.1 Summary of results

The odor tuning of PNs and KCs is compared, by means of field, tetrode, patch clamp and sharp electrode recordings, as well as pharmacological manipulation, in a locust *in vivo* preparation. Odor responses are broadly distributed over the PN population, with individual PNs' responses consisting of bouts of excitation and inhibition that outlast the stimulus by several seconds. A given PN's probability

of response to an odor is ~0.6 and the number of spikes fired ~20. This is in sharp contrast to the responses of the PNs' downstream decoders in the mushroom body (MB), the KCs, which respond to odors very sparsely (p(resp) ~0.1) and very briefly (~2 spikes). Two different mechanisms are identified that contribute to this transformation: an intrinsic biophysical property of the KCs, and a network effect resulting from oscillatory feedforward inhibition. In addition to forming synapses onto the KCs, PNs also contact inhibitory neurons in the lateral horn (LHIs), which in turn feed back onto the KCs. As a result of PN locking to the field oscillation, IPSPs from the LHIs arrive at the KCs systematically after the PN EPSPs, and thus prohibit unsynchronized PN spikes from influencing KC firing. This effectively tunes the KCs to synchronized PNs, a property which is further enhanced by an intrinsic voltage-dependence which nonlinearly amplifies coincident EPSPs.

4.1.2 Significance of results

These results are of relevance to a number of issues related to neural coding, including the role of oscillations, sparse coding, and the piecewise decoding of densely distributed representations. The combination of intrinsic biophysical properties and a ubiquitous circuit element, feedforward inhibition, is of general interest as a mechanism for generating a population of coincidence detectors. From a sensory systems perspective, the dense representation at the level of the antennal lobe appears useful for computing different aspects of the stimulus, while the sparse representation in the MB is likely to benefit memory formation, as it requires the manipulation of relatively few synapses, and minimizes overlap between different memories.

4.2 Spike-Timing Dependent Plasticity in the Olfactory System

Spike-timing dependent plasticity (STDP) refers to persistent changes in synaptic strength that result from paired pre- and postsynaptic activity on a millisecond timescale (Magee and Johnston, 1997; Markram et al., 1997; Bi and Poo, 1998). It has been described in many vertebrate brain structures and cell types, and is thought to underlie associative memory (Dan and Poo, 2006). The insect brain structure predominantly implicated in associative memory is the mushroom body (Strausfeld et al., 1998). Synapses onto MB extrinsic cells formed by KC axons in the MB lobes, (where neuromodulatory cells thought to mediate reward stimuli also project,) have been hypothesized to be involved in the memory storage of sparse KC patterns (Heisenberg, 2003).

4.2.1 Summary of results

The synaptic connections between KCs and MB extrinsic neurons (with dendritic fields predominantly in the β-lobe, βLNs) are studied by sharp electrode and field recordings, electrical stimulation and rudimentary computer modeling. Individual KCs can form strong synaptic contacts (mean ~1.5 mV), and synaptic weights from a given KC can differ depending on the postsynaptic βLNs. The strength of these connections can be modified by paired pre- and postsynaptic activity, characterized by a Hebbian STDP curve on a ±25 ms timescale. In the context of the odor-evoked 20-30 Hz oscillation generated in the AL and propagated through the MB, this should have a homeostatic fine-tuning effect on the phase of βLN spiking, which is found to be precise, as measured relative to the field oscillation, or relative to other βLNs recorded simultaneously. The synaptic fine-tuning mechanism is confirmed

by a simple computer model, and by direct experimental manipulation of odor-evoked βLN phase in a manner predicted by the STDP curve.

4.2.2 Significance of results

These results establish that STDP exists in invertebrates, in the MB, which is thought to play a role in associative memory. They also provide experimental evidence for the role of STDP as a mechanism for maintaining synchrony in activity propagating across multiple layers, as predicted previously by theoretical studies (Suri and Sejnowski, 2002; Zhigulin et al., 2003; Arthur and Boahen, 2006).

Appendix A

STDP in a β-lobe Network Model with Lateral Inhibitory Connections

The limited modeling data in Chapter 3 address how the experimentally derived STDP curve would be expected to influence the firing properties of neurons in the β-lobe, particularly the phase relative to the local field potential (LFP). These data are from a single-cell integrate-and-fire model, and represent the timecourse of one LFP cycle. In this appendix, preliminary results are presented from a network model simulating β-lobe activity over multiple LFP cycles. The motivation for considering such a network model derives from the experimental observation (included herein) that subsets of β-lobe neurons (βLNs) are synaptically connected to each other, constituting lateral inhibiton.

A.1 Introduction

Previous work has shown that for projection neurons (PNs) and for Kenyon cells (KCs) a single oscillation cycle is a meaningful unit for the encoding and decoding of olfactory information (Perez-Orive et al., 2002; Stopfer et al., 2003; Mazor and Laurent, 2005). Whether or not this is also the case for the KCs' down-

stream targets in the β-lobe was previously unknown. The question was addressed to some extent in Chapter 3, but a direct test of the timescale at which βLNs integrate their KC inputs was not included. The data in Figure A.1 were intended to answer this question. Multiple extracellular stimulating electrodes were embedded in different locations among the KC somata and stimulated at several intervals, concomitant with intracellular recording of a downstream βLN. The purpose of this experiment was to determine whether odor-evoked KC EPSPs are summated across multiple LFP cycles, or whether their precise temporal structure is preserved at this layer. When the stimulation interval was 0ms or 10ms, significant summation was observed (Figure A.1ai and ii), but at intervals of 30ms or more, there was none. The decay of the stimulation-evoked waveforms appeared much faster than observed in simultaneous dual intracellular recording of connected KC-βLN pairs (Figure 3.1, Chapter 3).

A.2 Results

A.2.1 Lateral inhibition in the β-Lobe in vivo

Figure A.1a iii suggests that such narrowing of EPSPs occurs at least partly as a result of feedforward inhibition. It shows the result of electrical stimulation of

---

Figure A.1 ***(on the next page).*** **Lateral Inhibition in the β-Lobe. a i**, Temporal integration for three different stimulation intervals, raw traces and averages. **a ii**, EPSP amplitudes for five different stimulation intervals. **a iii**, Three different responses (for three different stimulating electrode locations among KC somata) in a single postsynaptic βLN, raw traces and averages. **b i**, Example βLN and LFP recording illustrating odor-evoked EPSPs (grey arrows) and IPSPs (red arrows). **b ii**, Phase plot of odor evoked IPSPs, histograms (blue), mean IPSP phase (red) and mean βLN spike phase (AP). **b iii**, Two distinct examples of non-reciprocally connected βLN pairs. Traces shown are presynaptic-spike-triggered averages of postsynaptic Vm (black), and presynaptic-spike-triggered averages of presynaptic Vm to illustrate timing relationship (red).

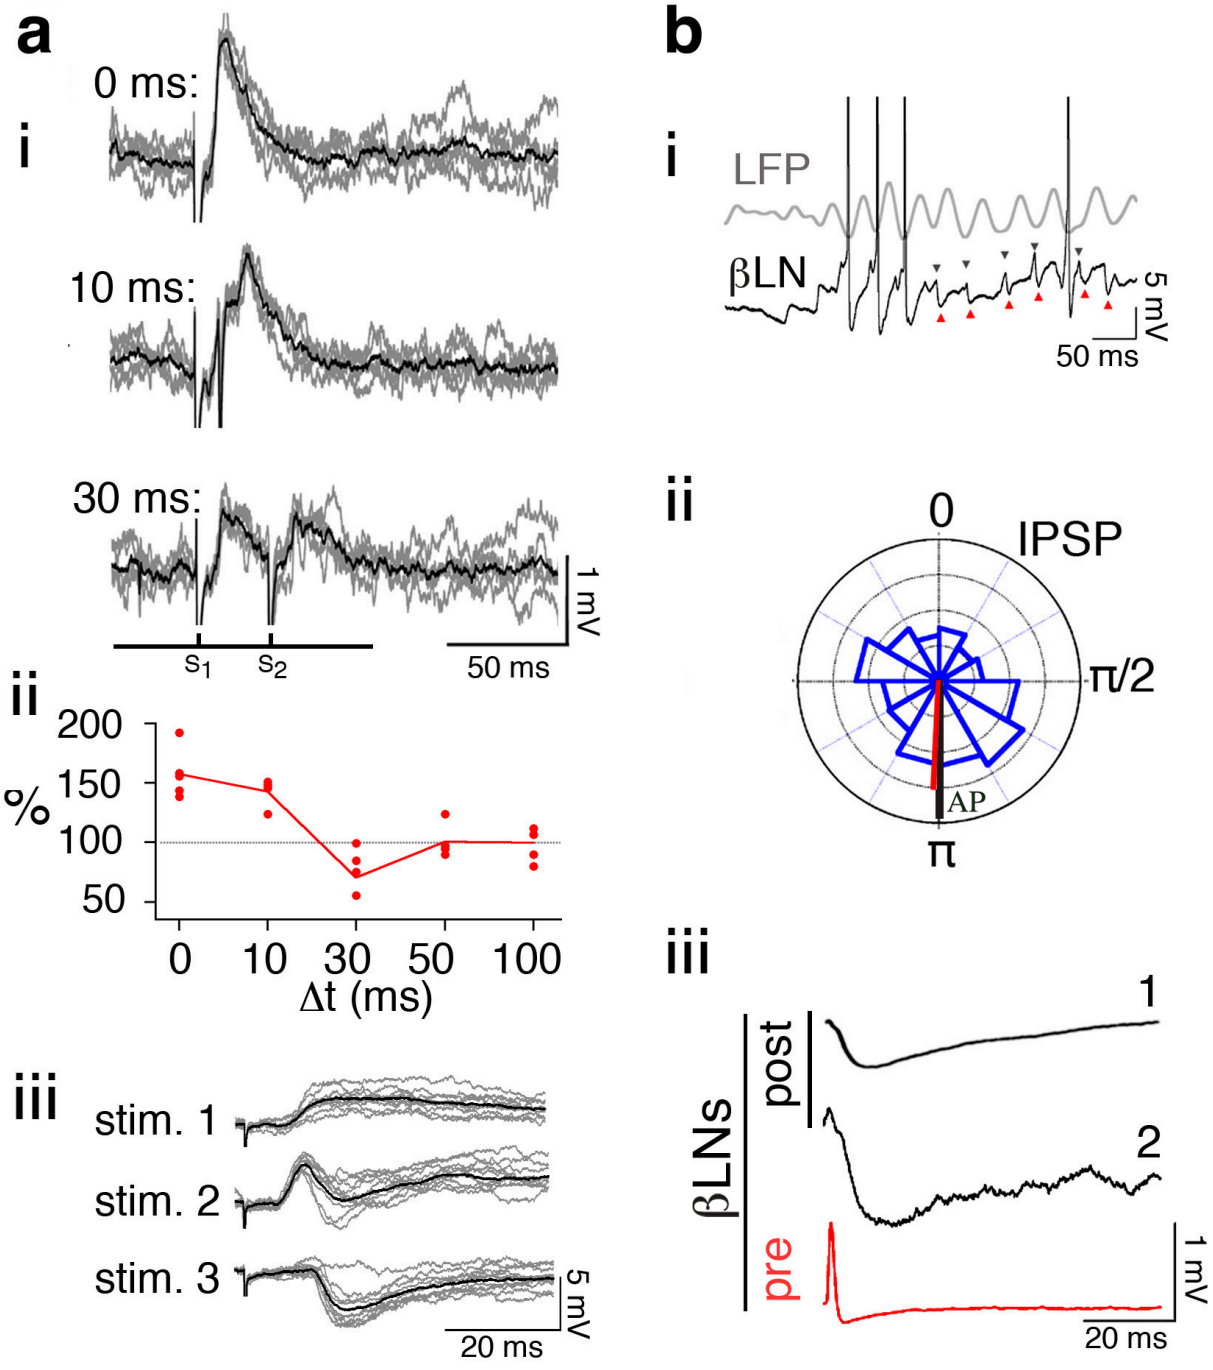
a
i
0 ms:
10 ms:
30 ms:
1 mV
S1
S2
50 ms
ii
200
150
100
50
%
0
10
30
50
100
Δt (ms)
iii
stim. 1
stim. 2
stim. 3
5 mV
20 ms
b
i
LFP
βLN
5 mV
50 ms
ii
0
IPSP
π/2
AP
π
iii
βLNs
post
pre
1
2
1 mV
20 ms

KCs at three different locations (stimulated at intervals of several seconds). The first electrode evoked a postsynaptic response with kinetics similar to the dual intracellular recordings (Figure 3.1, Chapter 3). The second electrode gave rise to a much sharper waveform, more similar to those in Figure A.1a i, while the third resulted in an IPSP, delayed by a few milliseconds. The onset of the IPSP evoked by the third electrode was well-alligned with the start of the decay in the second waveform, suggesting that the latter resulted from an EPSP followed by an IPSP. The interpretation is that stimulation by the second electrode activated KCs that were directly connected to the recorded βLN, as well KCs[1] connected to interposed inhibitory neurons. These neurons were brought to threshold by the activation of their KC inputs, causing them, in turn, to release inhibitory transmitter onto the recorded βLN. Presumably, the first electrode did not sufficiently activate inhibitory neurons connected to the recorded cell. In contrast, the third electrode may have stimulated a suitable subset of KCs for bringing interposed inhibitory neurons to threshold, while the recorded cell itself appeared not to be connected to any of these KCs. Such a succession of EPSPs and IPSPs was also observed in response to odor stimulation (Figure A.1bi and ii - grey and red arrows). Sometimes apparently pure IPSPs could be seen as well (Figure A.1bi, last red arrow).

We next sought to determine the identity of the cells that provide this feed-forward inhibition. The timing of the IPSPs[2] suggested that the βLNs themselves might be inhibiting each other. This is supported by the fact that the cells' somata are located within a cluster of inhibitory neurons in the Lateral Horn[3]. Direct ev-

[1] Possibly the same, or overlapping subset connected to the recorded βLN.

[2] i.e. immediately after the preferred βLN phase, as can be seen in the specific instance in Figure A.1bi, and on average in Figure A.1bii.

[3] As revealed by intracellular cobalt staining (as well as dye injection experiments by Sarah Farivar, 2005) and GABA immunohistochemistry by Glenn Turner (Perez-Orive et al., 2002).

idence came from simultaneous dual intracellular recording of connected βLNs. Two examples of non-reciprocally connected pairs (recorded in different animals) are shown in Figure A.1b iii. These connections can be very efficacious (as high as 5mV, data not shown).

These data answer the original question affirmatively, namely that cycle identity information is preserved across this synaptic layer, as it is between PNs and KCs. In addition, they reveal that the βLNs form a network connected by lateral inhibition. The IPSPs can be thought of as informing an individual βLN about the extent to which an odor is already represented at a particular LFP cycle. If there is no significant amount of inhibition at its typical phase (Figure A.2), and if it receives a moderate[4] level of KC activity, it will spike later than the typical βLN phase, π. This should result in net potentiation of its weights, and upon subsequent presentations, the cell will also fire at the typical phase. Given the data, which demonstrate very precise synchrony among subsets of βLNs, and tight locking to the LFP (Chapter 3), one might presume that the synchronous spikes are most relevant in terms of the βLN code. Under this assumption, the cell just described did not previously participate (significantly) in the encoding of the odor at the cycle in question when its spike occurred late. However, as a result of the potentiating effect of that same spike, the cell *will* subsequently become a member of the synchronous assembly.

If, however, the odor was already well represented in this cell's neighborhood (as defined by the βLNs synaptically connected to it), the cell would have received inhibition after the typical βLN phase, π, and the cell could not have produced its potentiation evoking spike. This kind of winner-take-all arrangement could be expected to level the network's activity across LFP cycles, to some extent, because

[4]where moderate means not enough to make it spike at π, but a sufficient level such that when summating over just a few 10s of milliseconds, threshold is reached in the potentiation window.

in cycles where inputs are strong, many cells fire, and weights should decrease due to the bias against potentiation effected by lateral inhibition. Conversely, where inputs are weaker, weights could be expected to grow more freely in the relative absence of inhibition. Of course, STDP, by itself already has the effect of setting the weights at a level that ensures an average phase of π, which implies activity at every cycle where the cell can reach threshold (in the potentiation window) at least occasionally. The main purpose of the modeling effort described below, is to assess how the interaction between STDP and lateral inhibition affects the activity patterns in the network.

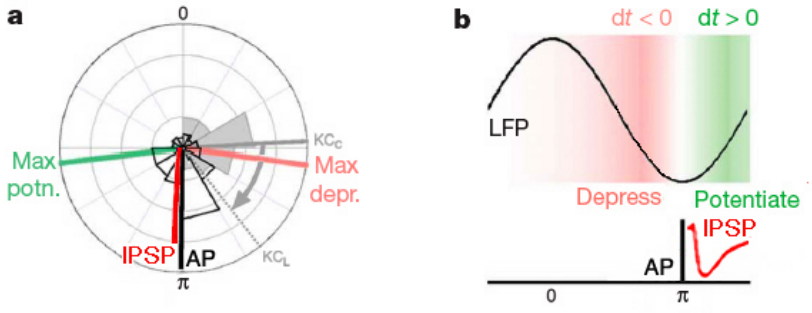

Figure A.2. **β-Lobe IPSP phase diagram. a**, Phase plot. KC spike time histograms (solid grey bins), KC mean spike phase as recorded near the calyx (KC_C), and when the KC spike arrives at the β-Lobe (KC_L; delay due to propagation along KC axons); βLN spike time histograms (black bins), βLN mean spike phase (AP); mean phase of βLN IPSP (IPSP); phases of maximum potentiation (Max potn.) and maximum depression (Max depr.). **b**, Diagram illustrating relative timing of βLN spike (AP), IPSP, and windows of depression and potentiation.

A.2.2 β-Lobe network model combining STDP and lateral inhibition

A network model of Izhikevich units (Izhikevich, 2007a) was used to simulate 30 βLNs[5]. Model parameters[6] are based on estimates from experimental data (Perez-

[5]Number estimated from diI-injections.

[6]A systematic exploration of parameter ranges has not yet been completed but the findings presented below appear to be quite robust against a relatively wide range of parameter values.

Orive et al., 2002; Mazor and Laurent, 2005; Cassenaer and Laurent, 2007). Excitatory inputs to the βLNs come from KCs (connectivity 10%, i.e. 30-50 KCs per βLN per LFP cycle); inhibitory inputs come from βLNs (connectivity 25%). KCs spike twice per odor presentation, and the specific firing times are drawn from gaussian distributions[7] (Figure A.3a). Parameters characterizing the membrane potential dynamics of individual units were chosen to give rise to single-cell subthreshold and spiking activity similar to experimental data (Figure A.3b), and were within the same range as in other studies (Izhikevich, 2007a). This resulted in network activity that appears similar to experimental data (Figure A.3c).

A.2.3 Effect of STDP and lateral inhibition on firing phase

The model reproduces the main result from the single-cell model in Chapter 3, namely that the experimentally derived STDP rule drives the average βLN firing phase to the trough of the LFP cycle, designated π (Figure A.4). The network converges to the same average phase, whether inhibition is implemented (Figure A.4a) or not (Figure A.4b and Figure 3.4e, Chapter 3). The phase at early and late cycles is initially delayed, because the drive from KCs is less than for intermediate

[7] Two distributions characterize the particular cycles at which KCs fire (one for each of the two spikes), and another determines the phase within a cycle.

Figure A.3 ***(on the next page).*** **Network model. a**, KC inputs to the β-Lobe network. **i**, Experimental KC activity profile, adapted from Mazor and Laurent (2005). Black bar indicates odor presentation. Blue box indicates time window used for simulations. **ii**, Profile approximating experimental data in (i), used to represent KC input onto β-LNs in the model; top blue curve is the sum of the lower two curves, which represent population PSTHs of model KCs' first (blue) and second (green) spikes, respectively. **iii**, Raster showing spike times of model KCs during a single trial. Each KC fires twice, in blue and green, respectively, as in histograms of (ii). **b**, Example voltage traces of β-LN responses. **i**, Experimental trace. **ii**, Model trace. **iii**, 10 consecutive trials for a single model β-LN. **c**, Rasters for 30 model β-LNs, 20 consecutive trials.

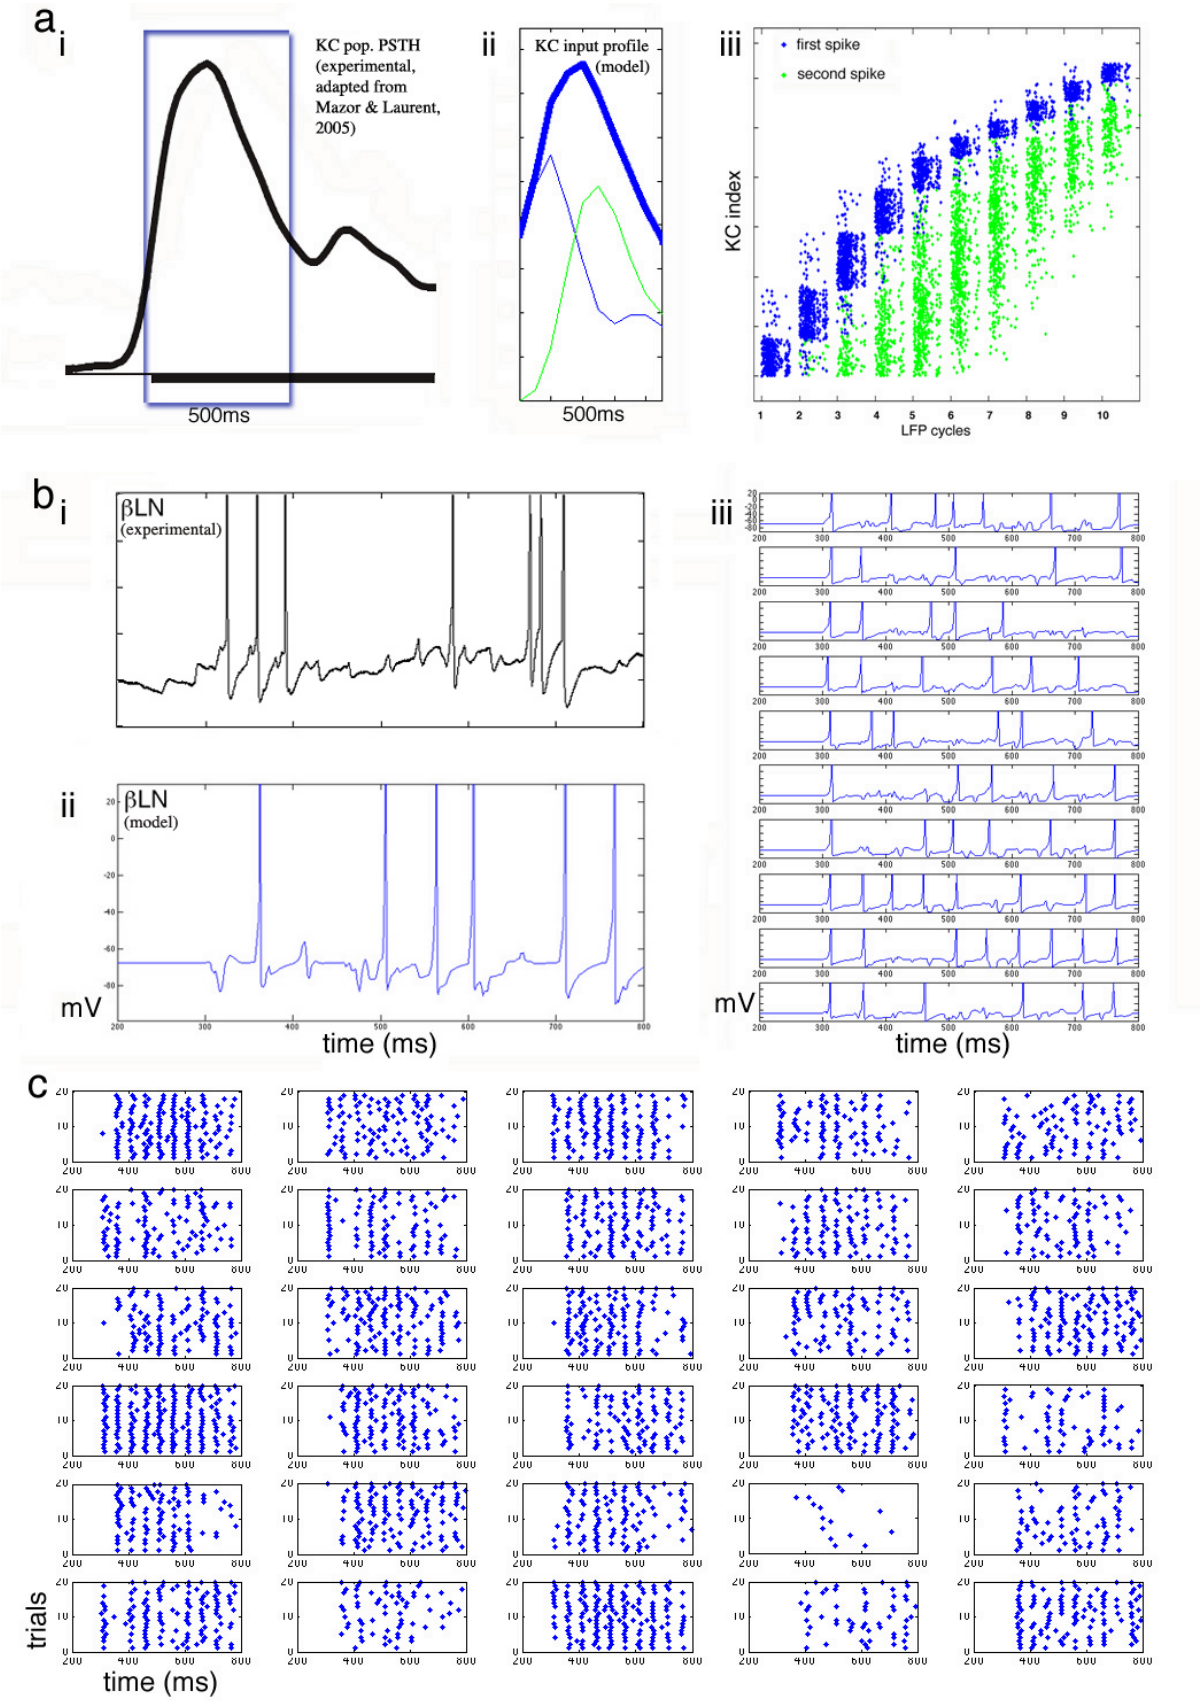
a
i
KC pop. PSTH
(experimental,
adapted from
Mazor & Laurent,
2005)
500ms
ii
KC input profile
(model)
500ms
iii
first spike
second spike
KC index
LFP cycles
b
i
βLN
(experimental)
ii
βLN
(model)
mV
time (ms)
iii
mV
time (ms)
c
trials
time (ms)

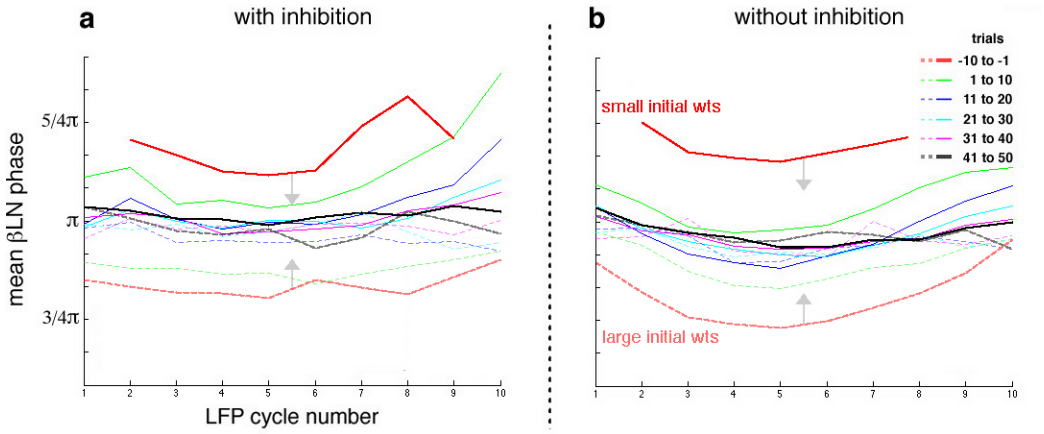

Figure A.4. The effect of STDP on βLN firing phase across multiple LFP cycles. a, With lateral inhibition. **b**, Without lateral inhibition. Results are shown for two starting weight distributions (initial weight distribution with small mean, solid curves; initial weight distribution with large mean, dashed curves). Curves are averages computed over 10 consecutive trials, as well as 10 distinct simulations with different seeds. Trials -10 to -1 are without STDP, trials 1 to 50 are with STDP. Grey arrows illustrate convergence from respective initial conditions.

cycles (Figure A.3a), but the STDP rule quickly (approximately) equalizes phase across all cycles. A similar evolution and convergence is seen, whether the initial KC-βLN weight distributions have a low (solid curves) or high mean (dashed curves).

A.2.4 Effect on average population activity

When inhibition among βLNs is implemented, activity across the population converges to a level such that, on average, approximately half (0.54) of the population is active per cycle (Figure A.5a, black curve). This is not the case, however, if lateral inhibition is eliminated from the model (Figure A.5b, average fraction of the population active: 0.85, black curve). In other words, the combined effect of lateral inhibition and STDP makes the network converge to an activity level approximating the middle of its dynamic range.

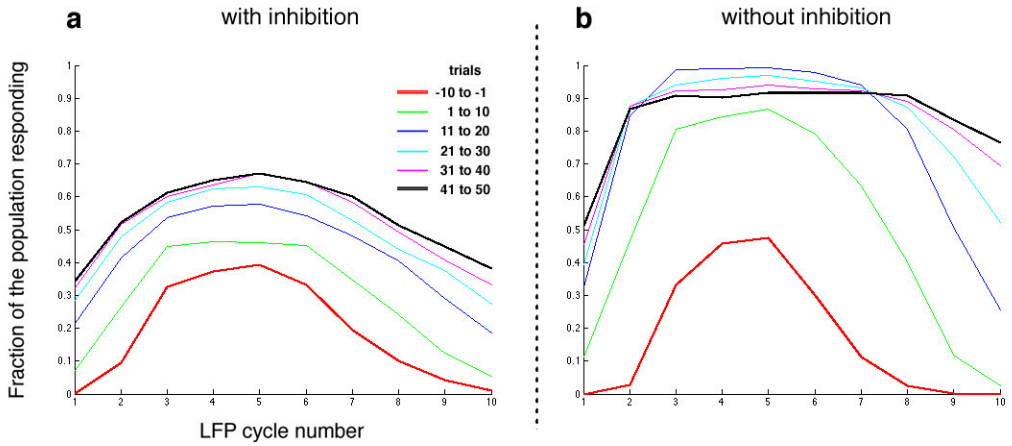

Figure A.5. The effect of STDP on βLN population activity across multiple LFP cycles. **a** With lateral inhibition. **b**, Without lateral inhibition. Results are shown for one starting weight distribution (with small mean). Curves are averages computed over 10 consecutive trials, as well as 10 distinct simulations with different seeds. Trials -10 to -1 are without STDP, trials 1 to 50 are with STDP.

A.2.5 Biased STDP

The fact that the network settles to the middle of its dynamic range under baseline conditions could be of particular relevance in the context of reinforcement learning. It is conceivable that plasticity at this synapse is affected by neuromodulators, such as would be released during reward or punishment. If, following Izhikevich (2007b), that interaction consists of a potentiation or depression bias of the STDP curve, then a default response with half of the network active should provide a large dynamic range to accommodate such contingencies. Figure A.6 illustrates this explicitly. When lateral inhibition is present, a 5-fold potentiation bias results in a 51% increase in the population's activity in just a few trials. The same bias for the same number of trials gives rise to an 18% increase when inhibition is absent, as the network has already converged closer to it's peak activity (due to the unconstrained effect of STDP). A 5-fold depression bias results in a 48% decrease in population activity for the inhibition-implementing network. The same bias evoked a 54%

decrease in the network without inhibition.

A.2.6 Reshaping of population activity profile

In contrast to the effect on phase (Figure A.4), the network's average activity level does not equalize across cycles[8] when lateral inhibition is present. Instead, the distribution of βLN activity across LFP cycles approximates the shape of the KC input distribution (Figure A.7a, compare solid and dashed black versus red curves; but not when inhibition is absent, Figure A.7b).

A.2.7 Multimodal distributions of initial weights and KC activity profiles

Figure A.7a demonstrates that the network reproduces its input distribution, given initial weight distributions that are *equal* across cycles[9]. Presumably, it should also be able to correct for weight distributions that are not homogeneous across cycles. This is in fact the case, as shown in Figure A.8a, where the means of the initial weight distributions corresponding to a subset of cycles (marked by *) are decreased by 25%, relative to the other cycles. The resultant βLN activity profiles (Figure A.8a, solid curves) are compared to those resulting from equal weight distributions as in Figure A.5a (reproduced as dashed curves in Figure A.8a). The initial response due to the unequal weight distributions is quite different, but application of STDP converges to a similar βLN activity distribution as in Figure A.5a, thus recovering the underlying KC activity profile.

The KC input distribution used so far is derived from data pooled across multiple KCs and multiple odors, resulting in a relatively smooth unimodal distribution.

[8] after 50 trials, as in Figure A.4.

[9] Weight distributions can be considered to be assigned to particular cycles, to the extent that the corresponding KCs fire reliably at those cycles.

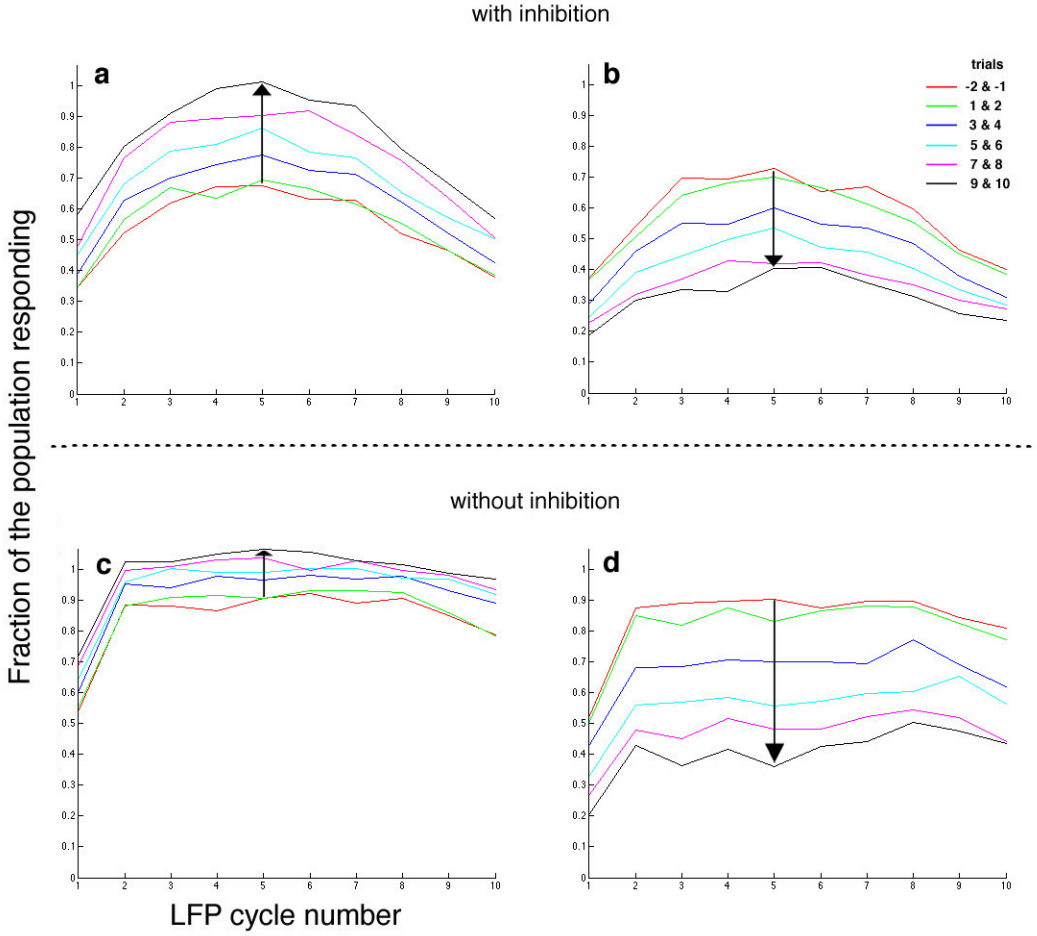

Figure A.6. The effect of STDP *bias* on βLN population activity across multiple LFP cycles. **a** and **b**, With lateral inhibition. **c** and **d**, Without lateral inhibition. Results are shown for starting weight distributions derived from from 50 trials of STDP (i.e. underlying the black curves in Figure A.5). Curves are averages computed over 2 consecutive trials (unlike the curves in other figures), as well as 10 distinct simulations with different seeds. Trials -2 to -1 are without STDP, trials 1 to 10 are with STDP. Arrows emphasize the effect of potentiation bias (**a, c**) and depression bias (**b, d**), respectively.

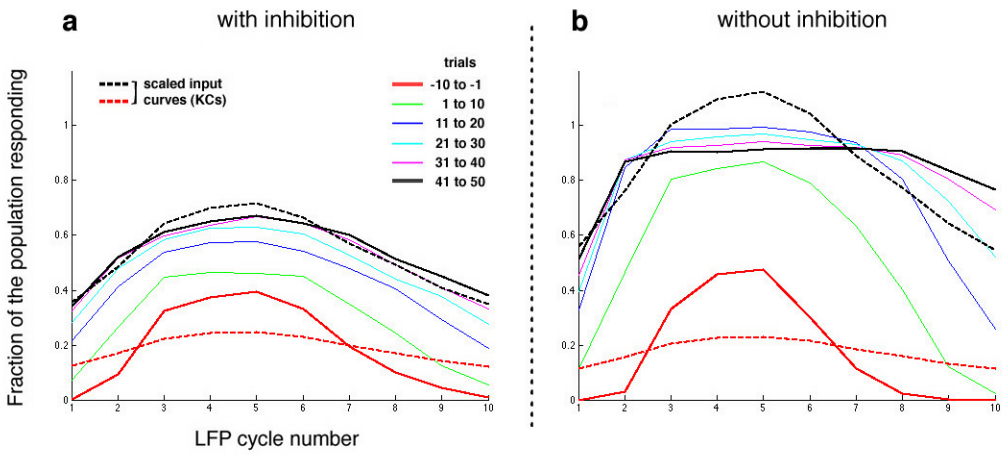

Figure A.7. Model β-Lobe reproduces KC activity profile across multiple LFP cycles. a With lateral inhibition. **b**, Without lateral inhibition. Results are shown for one starting weight distribution (with small mean). Curves are averages computed over 10 consecutive trials, as well as 10 distinct simulations with different seeds. Trials -10 to -1 are without STDP, trials 1 to 50 are with STDP. Solid curves are identical to those in Figure A.5. Dashed curves are superimposed KC input profiles, scaled to equal the integral of the initial (red) and final (black) βLN response curves, respectively.

As such, the change from the initial βLN activity profile[10] to the one more similar to the KC input profile (Figure A.7a) is perhaps not overly dramatic. However, when KC activity profiles are segregated by their respective odors, the distributions are not always unimodal[11]. The network's ability to closely reproduce such a single-odor KC activity profile is demonstrated in Figure A.8b.

A.3 Discussion

A reason for why this combination of features might be useful becomes apparent when considering the neural architecture in which this circuit is embedded, *in vivo*.

[10] which, in addition to the smooth KC input profile, furthermore derives from a normally distributed weight distribution.

[11] These differences appear to be significant, despite the relatively small number of KCs making up these profiles. The multimodal KC response profiles are also observed in a much larger KC data-set (Kai Shen, personal communication).

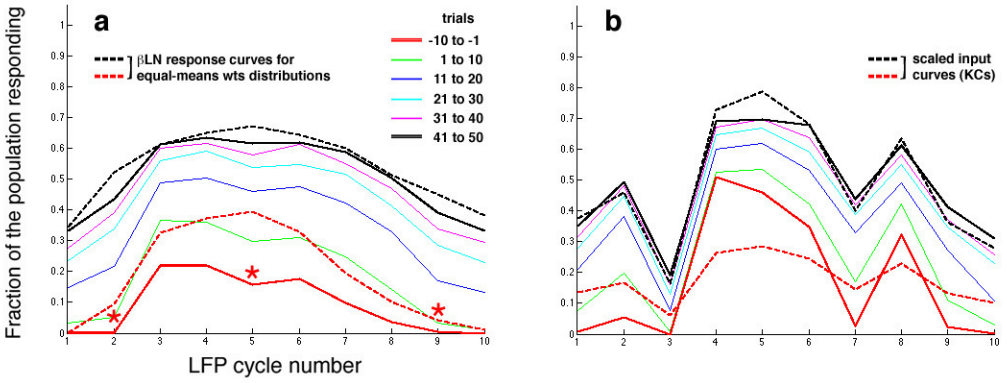

Figure A.8. Model β-Lobe response to multimodal synaptic weights distribution or input profile. a, Initial KC-βLN weights distribution not equal across cycles. All weights for KCs that fire on average during cycles 2, 5, and 9 (*) are decreased by 25% at onset of simulation. Curves are averages computed over 10 consecutive trials, as well as 10 distinct simulations with different seeds. Trials -10 to -1 are without STDP, trials 1 to 50 are with STDP. Dashed curves are responses to the same input profiles as the solid curves, but for equal initial weights distribution - reproduced from Figure A.5a. **b**, Multimodal KC input profile. KC input profile approximating experimental KC population PSTH for a *single* odor (based on data from Mazor and Laurent, 2005). Dashed curves are superimposed KC input profiles, scaled to equal the integral of the initial (red) and final (black) βLN response curves, respectively.

Projection neuron (PN) axons contact KC dendrites in the mushroom body (MB) calyx, and bifurcate, targeting the lateral horn (LH). βLNs also have projections in the LH, and, because of their dense arborizations among KC axons in the β-Lobe, are well suited to read out KC population activity and convey the result to the LH. If the β-Lobe provides inhibition that strongly resembles the KC input profile[12], then the LH could find its excitatory (PN) and inhibitory (βLN) input profiles well matched[13].

Summarizing these features leads to the following simplified model of olfactory learning. Prior to learning any association, the effect of odor-induced activation of

[12]Due to the combined effect of lateral inhibition and STDP within the β-Lobe
[13]To the extent that the global KC output well-represents its PN input.

PNs, kept in check by the output of the β-Lobe in the LH, would be a naive, intrinsic behavior[14]. When exposure to the odor is followed by an aversive or appetitive stimulus, an odor-specific set of KC-βLN synapses would be either strengthened or weakened, thus changing the gain on the PN drive into the LH, which would be hypothesized to alter the animal's odor-evoked behavior. This model predicts that the memory is stored at synapses between KCs and βLNs in manner that is not only odor- and KC-specific, but also also cycle-specific. This selectivity derives, indirectly, from the homeostatic effect of *un*biased STDP[15].

A few presentations of an odor followed by an unconditioned stimulus (U.S.)[16] should give rise to a specific strengthening (or weakening) of the synapses formed onto βLNs by KCs that were active during the odor presentation, as these would be the only KCs activated in immediate temporal proximity to postsynaptic spikes. The result would be β-lobe output onto the LH that exceeds (or falls short of) PN output onto the LH. Although the modification would be specific to the KC-βLN synapses activated by the odor, it would be expected to apply to all LFP cycles equally.

If the odor is presented several times afterwards *without* the U.S., the homeostatic effect of unbiased STDP should bring the synaptic weights that were previously modified back to baseline levels[17], which would be expected to underly a behaviorally measurable extinction of the memory. However, consider instead the scenario where the initial pairing is followed by a *different* odor presented without the U.S.. Any part of the PN trajectory that this odor would have in common with the previously paired odor would give rise to activation of a common set of

[14]This might be approach or withdrawal, or neither, depending on the particular odor.
[15]i.e. not followed by appetitive or aversive stimuli
[16]hypothesized to result in an STDP curve biased towards potentiation (or depression)
[17]i.e. synaptic weights that put the βLN back to the middle of its dynamic range

KCs. Re-activation of these KCs (without U.S.) should equally bring the KC- βLN synaptic weights back to equilibrium, much as if they had been activated by the first odor, as in the case of extinction. In other words, sequential exposure to multiple odors will sculpt the gain exerted by the MB, to the extent that odors with different contingencies evoke PN trajectories that traverse similar regions. Given that KCs tend to fire at particular LFP cycles in an odor- and KC-specific manner, the synaptic modification would effectively be cycle specific. And since the KCs' cycle-specificity corresponds to the reformatting across LFP cycles carried out by the antennal lobe (AL), particular odor *features* could be selectively rewarded. This would permit the animal to learn, for example, whether an entire odor class predicts a punishment (or reward), or whether instead only an individual member of the class is the reliable predictor.

Based on the above results, STDP is proposed to play multiple roles: to facilitate the synchronous flow of olfactory information and, more speculatively, to use its own fine-timescale sensitivity to select the specific set of synapses that should be modified if a neuromodulator were to be contingent on the odor. Also, in conjunction with lateral inhibition: to maintain activity in the β-Lobe near the middle of its dynamic range, and, lastly, to sculpt β-Lobe output so as to match its input profile across cycles of the synchronous oscillation.

References

Abarbanel HD, Huerta R, Rabinovich MI (2002) Dynamical model of long-term synaptic plasticity. *Proc Natl Acad Sci U S A* 99:10132–10137.

Abbott L, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101.

Abbott LF, Nelson SB (2000) Synaptic plasticity: taming the beast. *Nat Neurosci* 3 Suppl:1178–83.

Abeles M (1982) Role of the cortical neuron: integrator or coincidence dectector? *Isr. J. Med. Sci.* 18:83–92.

Abeles M, Bergman H, Margalit E, Vaadia E (1993) Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol.* 70:1629–38.

Abraham WC, Williams JM (2003) Properties and Mechanisms of LTP Maintenance. *Neuroscientist* 9:463–474.

Adrian E (1942) Olfactory reactions in the brain of a hedgehog. *J. Physiol. (Lond.)* 100:459–73.

Adrian ED (1950) The electrical activity of the olfactory bulb. *E.E.G. Clin. Neurophysiol.* 2:377–88.

Andersen P (2003) A prelude to long-term potentiation. *Phil Trans R Soc Lond B* 358:613–615.

Aroniadou-Anderjaska V, Ennis M, Shipley M (1999) Dendrodendritic recurrent excitation in mitral cells of the rat olfactory bulb. *J. Neurophysiol.* 82:489–94.

Aroniadou-Anderjaska V, Zhou FM, Priest CA, Ennis M, Shipley MT (2000) Tonic and Synaptically Evoked Presynaptic Inhibition of Sensory Input to the Rat Olfactory Bulb Via GABAB Heteroreceptors. *J Neurophysiol* 84:1194–1203.

Arthur J, Boahen K (2006) Learning in silicon: Timing is everything In Sholkopf B, Weiss Y, editors, *Advances in Neural Information Processing*, pp. 75–82. MIT Press.

Barlow HB (1969) Pattern recognition and the responses of sensory neurons. *Ann N Y Acad Sci* 156:872–81.

Baxter J, Bartlett P, Weaver L (2001) Experiments with infinite-horizon, policy gradient estimation. *Journal of Artificial Intelligence Research* 15:351–381.

Bell CC, Han VZ, Sugawara Y, Grant K (1997) Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387:278–281.

Bender VA, Bender KJ, Brasier DJ, Feldman DE (2006) Two Coincidence Detectors for Spike Timing-Dependent Plasticity in Somatosensory Cortex. *J. Neurosci.* 26:4166–4177.

Bi GQ, Poo MM (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci* 18:10464–72.

Bliss TVP, Collingridge GL (1993) A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361:31–39.

Bliss TVP, Lømo T (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol* 232:331–356.

Blum KI, Abbott LF (1996) A model of spatial map formation in the hippocampus of the rat. *Neural Comput* 8:85–93.

Boettiger CA, Doupe AJ (2001) Developmentally restricted synaptic plasticity in a songbird nucleus required for song learning. *Neuron* 31:809–818.

Bofill-i Petit A, Murray A (2004) ynchrony detection and amplification by silicon neurons with stdp synapses. *IEEE Trans Neural Netw* 15:1296–1304.

Bohte SM, Mozer MC (2007) Reducing the variability of neural responses: A computational theory of spike-timing-dependent plasticity. *Neural Computation* 19:371–403.

Bozza T, Kauer J (1997) Odorant response properties of convergent olfactory receptor neurons. *J. Neurosci.* 18:4560–9.

Bragin A, Jando G, Nadasdy Z, Hetke J, Wise K, Buzsaki G (1995) Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. *J. Neurosci.* 15:47–60.

Broome BM, Jayaraman V, Laurent G (2006) Encoding and decoding of overlapping odor sequences. *Neuron* 51:467–482.

Brown S, Joseph J, Stopfer M (2005) Encoding a temporally structured stimulus with a temporally structured neural representation. *Nat Neurosci.* 8:1568–1576.

Buck LB (1996) Information coding in the vertebrate olfactory system. *Annu. Rev. Neurosci.* 19.

Buck L, Axel R (1991) A novel multigene family may encode odorant receptors: a molecular basis for odor recognition. *Cell* 65:175–87.

Buonviso N, Chaput M, Berthommier F (1992) Temporal pattern analyses in pairs of neighboring mitral cells. *J. Neurophysiol.* 68:417–24.

Burrows M, Boeckh J, Esslen J (1982) Physiological and morphological properties of interneurons in the deutocerebrum of male cockroaches with responses to female pheromones. *J. Comp. Physiol.* A 145:447–57.

Cain W, Potts B (1996) Switch and bait: probing the discriminative basis of odor identification via recognition memory. *Chem Senses* 21:35–44.

Cang J, Isaacson JS (2003) In vivo whole-cell recording of odor-evoked synaptic transmission in the rat olfactory bulb. *J Neurosci* 23:4108–16.

Carlsson MA, Galizia CG, Hansson BS (2002) Spatial representation of odours in the antennal lobe of the moth spodoptera littoralis (lepidoptera: Noctuidae). *Chem Senses* 27:231–44.

Cassenaer S, Laurent G (2007) Hebbian stdp in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature* 448:709–713.

Christensen TA, Harrow I, Cuzzocrea C, Randolph P, Hildebrand J (1995) Distinct projections of two populations of olfactory receptor axons in the antennal lobe of the sphinx moth manduca sexta. *Chem Senses* 20:313–23.

Cinelli A, Kauer J (1992) Voltage-sensitive dyes and functional activity in the olfactory pathway. *Annu. Rev. Neurosci.* 15:321–51.

Clyne PJ, Certel SJ, de Bruyne M, Zaslavsky L, Johnson WA, Carlson JR (1999) The odor specificities of a subset of olfactory receptor neurons are governed by acj6, a pou-domain transcription factor. *Neuron* 22:339–47.

Collingridge GL, Bliss TVP (1987) Nmda receptors - their role in long-term potentiation. *Trends in Neurosciences* 10:288–293.

Contreras D, Destexhe A, Steriade M (1997) Intracellular and computational characterization of the intracortical inhibitory control of synchronized thalamic inputs in vivo. *J. Neurophysiol.* 78:335–30.

Couto A, Alenius M, Dickson BJ (2005) Molecular, anatomical, and functional organization of the drosophila olfactory system. *Current Biology* 15:1535–1547.

Cragg B, Hamlyn L (1955) Action potentials of the pyramidal neurons in the hippocampus of the rabbit. *J Physiol* 129:608–627.

Csibra G, Davis G, Spratling MW, Johnson MH (2000) Gamma oscillations and object processing in the infant brain. *Science* 290:1582–5.

Dan Y, Poo MM (2006) Spike timing-dependent plasticity: from synapse to perception. *Physiol Rev* 86:1033–48.

Dan Y, Poo Mm (2004) Spike timing-dependent plasticity of neural circuits. *Neuron* 44:23–30.

Davis RL (2005) Olfactory memory formation in drosophila: From molecular to systems neuroscience. *Annual Review of Neuroscience* 28:275–302.

Davison IG, Katz LC (2007) Sparse and Selective Odor Coding by Mitral/Tufted Neurons in the Main Olfactory Bulb. *J. Neurosci.* 27:2091–2101.

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

de Bruyne M, Foster K, Carlson JR (2001) Odor coding in the drosophila antenna. *Neuron* 30:537–552.

Debanne D, Gahwiler BH, Thompson SM (1998) Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J Physiol* 507:237–247.

Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–33.

Drew PJ, Abbott LF (2006) Extending the effects of spike-timing-dependent plasticity to behavioral timescales. *Proc Natl Acad Sci U S A* 103:8876–81.

Dubnau J, Grady L, Kitamoto T, Tully T (2001) Disruption of neurotransmission in drosophila mushroom body blocks retrieval but not acquisition of memory. *Nature* 411:476–80.

Duchamp-Viret P, Chaput MA, Duchamp A (1999) Odor Response Properties of Rat Olfactory Receptor Neurons. *Science* 284:2171–2174.

Duchamp-Viret P, Duchamp A, Chaput MA (2000) Peripheral Odor Coding in the Rat and Frog: Quality and Intensity Specification. *J. Neurosci.* 20:2383–2390.

Dujardin F (1850) Memoire sur le systeme nerveux des insectes. *Ann. Sci. Nat. Zool.* 14:195–206.

Dujardin F (1853) Quelques observations sur les abeilles et particulierement sur les actes qui, chez ces insectes peuvent etre rapportes a l'intelligence. *Ann. Sci. Nat. Zool.* 18:231–240.

Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biol Cybern* 60:121–130.

Egger V, Feldmeyer D, Sakmann B (1999) Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nat Neurosci* 2:1098–1105.

Emptage NJ, Reid CA, Fine A, Bliss TVP (2003) Optical quantal analysis reveals a presynaptic component of ltp at hippocampal schaffer-associational synapses. *Neuron* 38:797–804.

Engel AK, Fries P, Singer W (2001) Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci* 2:704–16.

Ennis M, Zhou FM, Ciombor KJ, Aroniadou-Anderjaska V, Hayar A, Borrelli E, Zimmer LA, Margolis F, Shipley MT (2001) Dopamine D2 Receptor-Mediated Presynaptic Inhibition of Olfactory Nerve Terminals. *J Neurophysiol* 86:2986–2997.

Fahrbach SE (2006) Structure of the mushroom bodies of the insect brain. *Annu. Rev. Entomol.* 51:209–232.

Farivar SS (2005) Cytoarchitecture of the locust olfactory system Ph.D. diss., California Institute of Technology.

Feldman DE (2000) Timing-based ltp and ltd at vertical inputs to layer ii/iii pyramidal cells in rat barrel cortex. *Neuron* 27:45–56.

Felsen G, Shen Ys, Yao H, Spor G, Li C, Dan Y (2002) Dynamic modification of cortical orientation tuning mediated by recurrent connections. *Neuron* 36:945–954.

Florian RV (2007) Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. *Neural Computation* 19:1468–1502.

Frey U, Morris RGM (1998) Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends in Neurosciences* 21:181–188.

Fricker DM R (2000) Epsp amplification and the precision of spike timing in hippocampal neurons. *Neuron* 28:559–69.

Friedman D, Strowbridge BW (2003) Both electrical and chemical synapses mediate fast network oscillations in the olfactory bulb. *J Neurophysiol* 89:2601–10.

Friedrich RW, Habermann CJ, Laurent G (2004) Multiplexing using synchrony in the zebrafish olfactory bulb. *Nat Neurosci* 7:862–71.

Friedrich RW, Korsching SI (1997) Combinatorial and chemotopic odorant coding in the zebrafish olfactory bulb visualized by optical imaging. *Neuron* 18:737–52.

Friedrich RW, Laurent G (2001) Dynamic optimization of odor representations by slow temporal patterning of mitral cell activity. *Science* 291:889–94.

Friedrich RW, Laurent G (2004) Dynamics of olfactory bulb input and output activity during odor stimulation in zebrafish. *J Neurophysiol* 91:2658–69.

Fries P, Reynolds JH, Rorie AE, Desimone R (2000) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–63.

Froemke RC, Poo Mm, Dan Y (2005) Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* 434:221–225.

Galarreta M, Hestrin S (2001) Spike transmission and synchrony detection in networks of gabaergic interneurons. *Science* 292:2295–9.

Gao Q, Yuan B, Chess A (2000) Convergent projections of drosophila olfactory neurons to specific glomeruli in the antennal lobe. *Nat Neurosci* 3:780–85.

Gelperin A, Tank D (1990) Odour-modulated collective network oscillations of olfactory interneurons in a terrestrial mollusc. *Nature* 345:437–40.

Gerstner W (2001) A framework for spiking neuron models: The spike response model In Moss F, Gielen S, editors, *The Handbook of biological physics*, Vol. 4, pp. 469–516. Elsevier, Amsterdam.

Goldman AL, Van der Goes van Naters W, Lessing D, Warr CG, Carlson JR (2005) Coexpression of two functional odor receptors in one neuron. *Neuron* 45:661–666.

Gray C (1994) Synchronous oscillations in neuronal systems: mechanisms and function. *J. Comput. Neurosci.* 1:11–38.

Gray CM, Konig P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338:334–337.

Guyonneau R, VanRullen R, Thorpe SJ (2005) Neurons tune to the earliest spikes through stdp. *Neural Computation* 17:859–879.

Haberly L (1990) Comparative aspects of olfactory cortex. In E.G. Jones AP, editor, *Cerebral Cortex*, pp. 137–66. Plenum Press, New York.

Hallem EA, Ho MG, Carlson JR (2004) The molecular basis of odor coding in the drosophila antenna. *Cell* 117:965–79.

Hamilton KA, Kauer JS (1989) Patterns of intracellular potentials in salamander mitral,tufted cells in response to odor stimulation. *J Neurophysiol* 62:602–25.

Hansson BS, Anton S (2000) Function and morphology of the antennal lobe: new developments. *Annu Rev Entomol* 45:203–31.

Hansson BS, Carlsson M, B. K (2003) Olfactory activation patterns in the antennal lobe of the sphinx moth, manduca sexta. *J Comp Physiol A Neuroethol Sens Neural Behav Physiol.* 189:301–308.

Hayar A, Karnup S, Ennis M, Shipley MT (2004a) External tufted cells: a major excitatory element that coordinates glomerular activity. *J Neurosci* 24:6676–85.

Hayar A, Karnup S, Shipley MT, Ennis M (2004b) Olfactory bulb glomeruli: external tufted cells intrinsically burst at theta frequency and are entrained by patterned olfactory input. *J Neurosci* 24:1190–9.

Hebb D (1949) *The Organization of Behavior; a Neuropsychological Theory* Wiley, New York.

Heisenberg M (2003) Mushroom body memoir: from maps to models. *Nat Rev Neurosci* 4:266–75.

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

Heisenberg M (1980) Mutations of brain structure and function: What is the significance of the mushroom bodies for behavior? In Siddiqi O, P. B, Hall L, Hall

J, editors, *Development and Neurobiology of Drosophila*. Plenum Press, New York, NY.

HoneybeeGenomeConsortium (2006) Insights into social insects from the genome of the honeybee apis mellifera. *Nature* 443:931–49.

Hopfield JJ (1995) Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376:33–6.

Isaacson JS, Vitten H (2003) GABAB Receptors Inhibit Dendrodendritic Transmission in the Rat Olfactory Bulb. *J. Neurosci.* 23:2032–2039.

Isaacson J, Strowbridge B (1998) Olfactory reciprocal synapses: dendritic signaling in the cns. *Neuron* 20:749–62.

Izhikevich E (2007a) *Dynamical systems in neuroscience: the geometry of excitability and bursting* MIT Press.

Izhikevich EM (2007b) Solving the distal reward problem through linkage of stdp and dopamine signaling. *Cerebral Cortex* .

Jawlowski H (1958) Nerve tracks in bee (apis mellifera) running from the sight and antennal organs to the brain. *Ann. Univ. Marie Curie-Sklodowska* 12:307–323.

Jawlowski H (1960) On the brain structure of the symphyta (hymenoptera). *Bull. Acad. Polon. Sci.* 8:265–268.

Joerges j, kuettner a, Galizia CG, Menzel R (1997) Representations of odours and odour mixtures visualized in the honeybee brain. *Nature* 387:285–288.

Jortner R, Farivar SS, Laurent G (2007) A simple connectivity scheme for sparse coding in an olfactory system. *J Neurosci* 27:1659–1669.

Karmarkar UR, Buonomano DV (2002) A Model of Spike-Timing Dependent Plasticity: One or Two Coincidence Detectors? *J Neurophysiol* 88:507–513.

Karmarkar U, Najarian M, Buonomano D (2002) Mechanisms and significance of spike-timing dependent plasticity. *Biological cybernetics* 87:373–382.

Kempter R, Gerstner W, Hemmen JLv (2001) Intrinsic stabilization of output rates by spike-based hebbian learning. *Neural Computation* 13:2709–2741.

Kenyon F (1896) The brain of the bee. *J. Comp. Neurol.* 6:133–210.

Kistler WM (2002) Spike-timing dependent synaptic plasticity: a phenomenological framework. *BIOLOGICAL CYBERNETICS* 87:416.

Kistler WM, van Hemmen J (2000) Modeling synaptic plasticity in conjuction with the timing of pre- and postsynaptic action potentials. *Neural Computation* 12:385–405.

Kitajima T, Hara K (2000) A generalized hebbian rule for activity-dependent synaptic modifications. *Neural Networks* 13:445–454.

Koester H, Sakmann B (1998) Calcium dynamics in single spines during coincident pre- and postsynaptic activity depend on relative timing of back-propagating action potentials and subthreshold excitatory postsynaptic potentials. *Proc Natl Acad Sci U S A.* 95:9596–601.

Konig P, Engel AK, Singer W (1996) Integrator of coincidence detector? the role of the coritcal neuron revisited. *Trends Neurosci* 19:130–37.

Kubota S, Kitajima T (2007) A model for synaptic development regulated by nmda receptor subunit expression. *J Comput Neurosci* DOI:10.1007/s10827-007-0036-8.

Lam YW, Cohen LB, Wachowiak M, Zochowski MR (2000) Odors elicit three different oscillations in the turtle olfactory bulb. *J Neurosci* 20:749–62.

Laurent G (2002) Olfactory network dynamics and the coding of multidimensional signals. *Nat Rev Neurosci* 3:884–95.

Laurent G, Davidowitz H (1994) Encoding of olfactory information with oscillating neural assemblies. *Science* 265:1872–5.

Laurent G, Naraghi M (1994) Odorant-induced oscillations in the mushroom bodies of the locust. *J Neurosci* 14:2993–3004.

Laurent G, Stopfer M, Friedrich RW, Rabinovich MI, Volkovskii A, Abarbanel HD (2001) Odor encoding as an active, dynamical process: experiments, computation, and theory. *Annu Rev Neurosci* 24:263–97.

Laurent G, Wehr M, Davidowitz H (1996) Temporal representations of odors in an olfactory network. *J Neurosci* 16:3837–47.

Lazar A, Pipa G, Triesch J (2007) Fading memory and time series prediction in recurrent networks with different forms of plasticity. *Neural Networks* 20:312–322.

Lehmkuhle MJ, Normann RA, Maynard EM (2006) Trial-by-Trial Discrimination of Three Enantiomer Pairs by Neural Ensembles in Mammalian Olfactory Bulb. *J Neurophysiol* 95:1369–1379.

Leitch B, Laurent G (1996) Gabaergic synapses in the antennal lobe and mushroom body of the locust olfactory system. *J Comp Neurol* 372:487–514.

Leveteau J, MacLeod P (1966) Olfactory discrimination in the rabbit olfactory glomerulus. *Science* 153:175–6.

Li Cy, Lu Jt, Wu Cp, Duan Sm, Poo Mm (2004) Bidirectional modification of presynaptic neuronal excitability accompanying spike timing-dependent synaptic plasticity. *Neuron* 41:257–268.

Lin YW, Min MY, Chiu TH, Yang HW (2003) Enhancement of Associative Long-Term Potentiation by Activation of beta-Adrenergic Receptors at CA1 Synapses in Rat Hippocampal Slices. *J. Neurosci.* 23:4173–4181.

Linster C, Smith B (1999) Generalization between binary odor mixtures and their components in the rat. *Physiol. Behav.* 66:701–7.

Lisman J (1989) A Mechanism for the Hebb and the Anti-Hebb Processes Underlying Learning and Memory. *PNAS* 86:9574–9578.

Liu X, Davis RL (2006) Insect olfactory memory in time and space. *Current Opinion in Neurobiology* 16:679–685.

Livermore A, Laing D (1996) Influence of training and experience on the perception of multicomponent odor mixtures. *J Exp Psychol.* 22:267–77.

Lowe G (2003) Electrical signaling in the olfactory bulb. *Current Opinion in Neurobiology* 13:476–481.

Luo M, Katz LC (2001) Response correlation maps of neurons in the mammalian olfactory bulb. *Neuron* 32:1165–79.

MacLeod K, Backer A, Laurent G (1998) Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* 395:693–8.

MacLeod K, Laurent G (1996) Distinct mechanisms for synchronization and temporal patterning of odor-encoding neural assemblies. *Science* 274:976–9.

Macrides F, Chorover S (1972) Olfactory bulb units: activity correlated with inhalation cycles and odor quality. *Science* 185:84–7.

Magee JC, Johnston D (1997) A Synaptically Controlled, Associative Signal for Hebbian Plasticity in Hippocampal Neurons. *Science* 275:209–213.

Malenka RC, Bear MF (2004) Ltp and ltd: An embarrassment of riches. *Neuron* 44:5–21.

Malenka RC, Nicoll RA (1999) Long-Term Potentiation–A Decade of Progress? *Science* 285:1870–1874.

Malinow R, Malenka RC (2002) Ampa receptor trafficking and synaptic plasticity. *Annual Review of Neuroscience* 25:103–126.

Malnic B, Hirono J, Sato T, Buck LB (1999) Combinatorial receptor codes for odors. *Cell* 96:713–23.

Margrie TW, Sakmann B, Urban NN (2001) Action potential propagation in mitral cell lateral dendrites is decremental and controls recurrent and lateral inhibition in the mammalian olfactory bulb. *Proc Natl Acad Sci U S A* 98:319–24.

Margrie TW, Schaefer AT (2003) Theta oscillation coupled spike latencies yield computational vigour in a mammalian sensory system. *J Physiol* 546:363–374.

Markram H, Lubke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science* 275:213–5.

Masuda N, Aihara K (2004) Self-organizing dual coding based on spike-time-dependent plasticity. *Neural Computation* 16:627–663.

Masuda N, Aihara K (2007) Dual coding hypotheses for neural information representation. *Mathematical Biosciences* 207:312–321.

Mazor O, Laurent G (2005) Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* 48:661–73.

McGuire SE, Le PT, Davis RL (2001) The role of drosophila mushroom body signaling in olfactory memory. *Science* 293:1330–3.

McGuire SE, Deshazer M, Davis RL (2005) Thirty years of olfactory learning and memory research in drosophila melanogaster. *Progress in Neurobiology* 76:328–347.

Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* 417:741–6.

Meister M, Bonhoeffer T (2001) Tuning and topography in an odor map on the rat olfactory bulb. *J Neurosci* 21:1351–60.

Menzel R, Erber J, Masuhr T (1974) Learning and memory in the honeybee In Barton Browne L, editor, *Experimental Analysis of Insect Behaviour*. Springer, Heidleberg, Germany.

Meredith M (1986) Patterned response to odor in mammalian olfactory bulb: the influence of intensity. *J. Neurophysiol.*, 56.

Mombaerts P (2004a) Genes and ligands for odorant, vomeronasal and taste receptors. *Nat Rev Neurosci* 5:263–78.

Mombaerts P (2004b) Odorant receptor gene choice in olfactory sensory neurons: the one receptor-one neuron hypothesis revisited. *Curr Opin Neurobiol* 14:31–6.

Mombaerts P, Wang F, Dulac C, Chao SK, Nemes A, Mendelsohn M, Edmondson J, Axel R (1996) Visualizing an olfactory sensory map. *Cell* 87:675–86.

Morrison A, Aertsen A, Diesmann M (2007) Spike-timing-dependent plasticity in balanced random networks. *Neural Computation* 19:1437–1467.

Motokizawa F (1996) Odor representation and discrimination in mitral,tufted cells of the rat olfactory bulb. *Exp. Brain Res.* 112:24–34.

Nevian T, Sakmann B (2006) Spine Ca2+ Signaling in Spike-Timing-Dependent Plasticity. *J. Neurosci.* 26:11001–11013.

Nishiyama M, Hong K, Mikoshiba K, Poo Mm, Kato K (2000) Calcium stores regulate the polarity and input specificity of synaptic modification. *Nature* 408:584–588.

Nowotny T, Rabinovich MI, Huerta R, Abarbanel HD (2003) Decoding temporal information through slow lateral excitation in the olfactory system of insects. *J Comput Neurosci* 15:271–81.

Olsen SR, Bhandawat V, Wilson RI (2007) Excitatory interactions between olfactory processing channels in the drosophila antennal lobe. *Neuron* 54:89–103.

Patel AD, Balaban E (2000) Temporal patterns of human cortical activity reflect tone sequence structure. *Nature* 404:80–84.

Perez-Orive J, Bazhenov M, Laurent G (2004) Intrinsic and circuit properties favor coincidence detection for decoding oscillatory input. *J Neurosci* 24:6037–47.

Perez-Orive J, Mazor O, Turner GC, Cassenaer S, Wilson RI, Laurent G (2002) Oscillations and sparsening of odor representations in the mushroom body. *Science* 297:359–65.

Pfister JP, Toyoizumi T, Barber D, Gerstner W (2006) Optimal spike-timing-dependent plasticity for precise action potential firing in supervised learning. *Neural Computation* 18:1318–1348.

Pittenger C, Kandel ER (2003) In search of general mechanisms for long-lasting plasticity: Aplysia and the hippocampus. *Philos Trans R Soc Lond B Biol Sci* 358:757–763.

Pouille F, Scanziani M (2001) Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science* 293:1159–63.

Pouzat C, Mazor O, Laurent G (2002) Using noise signature to optimize spike-sorting and to assess neuronal classification quality. *J Neurosci Methods* 122:43–57.

Quinn WG, Harris WA, Benzer S (1974) Conditioned behavior in drosophila melanogaster. *Proc Natl Acad Sci U S A* 71:708–12.

Ressler K, Sullivan S, Buck L (1993) A zonal organization of odorant receptor gene expression in the olfactory epithelium. *Cell* 73:597–609.

Riehle A, Grün S, Diesmann M, Aertsen A (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 281:34–42.

Rinberg D, Koulakov A, Gelperin A (2006) Sparse Odor Coding in Awake Behaving Mice. *J. Neurosci.* 26:8857–8865.

Roberts PD, Bell CC (2002) Spike timing dependent synaptic plasticity in biological systems. *Biol Cybern* 87:392–403.

Rodriguez R, George M, Lachaux JP, Martinerie J, Renault B, Varela FJ (1999) Perception's shadow: long-distance synchronization of human brain activity. *Nature* 397:430–433.

Rolls ET, Tovee MJ (1995) Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J Neurophysiol* 73:713–26.

Rubin BD, Katz LC (1999) Optical imaging of odorant representations in the mammalian olfactory bulb. *Neuron* 23:499–511.

Rubin JE, Gerkin RC, Bi GQ, Chow CC (2005) Calcium Time Course as a Signal for Spike-Timing-Dependent Plasticity. *J Neurophysiol* 93:2600–2613.

Sattelle D, Buckingham S (2006) Invertebrate studies and their ongoing contributions to neuroscience. *Invertebrate Neuroscience* 6:1–3.

Schaefer AT, Margrie TW (2007) Spatiotemporal representations in the olfactory system. *Trends in Neurosciences* 30:92–100.

Schafer S, Rosenbloom J, Menzel R (1994) Ionic currents of kenyon cells from the mushroom body of the honeybee. *J. Neurosci.* 14:4600–12.

Schuett S, Bonhoeffer T, Hubener M (2001) Pairing-induced changes of orientation maps in cat visual cortex. *Neuron* 32:325–337.

Shadlen MN, Movshon JA (1999) Synchrony unbound: A critical evaluation of the temporal binding hypothesis. *Neuron* 24:67–77.

Shang Y, Claridge-Chang A, Sjulson L, Pypaert M, Miesenbock G (2007) Excitatory local circuits and their implications for olfactory processing in the fly antennal lobe. *Cell* 128:601–612.

Shepherd G (2004) *The Synaptic Organization of the Brain* Oxford University Press.

Shouval HZ, Bear MF, Cooper LN (2002) A unified model of nmda receptor-dependent bidirectional synaptic plasticity. *Proc Natl Acad Sci U S A* 99:10831–10836.

Shouval HZ, Kalantzis G (2005) Stochastic Properties of Synaptic Transmission Affect the Shape of Spike Time-Dependent Plasticity Curves. *J Neurophysiol* 93:1069–1073.

Sjostrom PJ, Turrigiano GG, Nelson SB (2001) Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32:1149–64.

Song I, Huganir RL (2002) Regulation of ampa receptors during synaptic plasticity. *Trends in Neurosciences* 25:578–588.

Song S, Miller KD, Abbott LF (2000) Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nature neuroscience* 3:919–926.

Spors H, Grinvald A (2002) Spatio-temporal dynamics of odor representations in the mammalian olfactory bulb. *Neuron* 34:301–15.

Spors H, Wachowiak M, Cohen LB, Friedrich RW (2006) Temporal Dynamics and Latency Patterns of Receptor Neuron Input to the Olfactory Bulb. *J. Neurosci.* 26:1247–1259.

Steriade M, Amzica F, Contreras D (1996) Synchronization of the fast (30-40hz) spontaneous cortical rhythms during brain activatio. *J. Neurosci.* 16:392–417.

Stewart W, Kauer J, Shepherd G (1979) Functional organization of rat olfactory bulb analysed by the 2-deoxyglucose method. *J. Comp. Neurol.* 185:715–34.

Stocker RF, Lienhard MC, Borst A, Fischbach KF (1990) Neuronal architecture of the antennal lobe in drosophila melanogaster. *Cell Tissue Res* 262:9–34.

Stopfer M, Bhagavan S, Smith BH, Laurent G (1997) Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 390:70–4.

Stopfer M, Jayaraman V, Laurent G (2003) Intensity versus identity coding in an olfactory system. *Neuron* 39:991–1004.

Stopfer M, Laurent G (1999) Short-term memory in olfactory network dynamics. *Nature* 402:664–8.

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

Stuart GJ, Sakmann B (1994) Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* 367:69–72.

Suh GS, Wong AM, Hergarden AC, Wang JW, Simon AF, Benzer S, Axel R, Anderson DJ (2004) A single population of olfactory sensory neurons mediates an innate avoidance behaviour in drosophila. *Nature* .

Suri RE, Sejnowski TJ (2002) Spike propagation synchronized by temporally asymmetric hebbian learning. *Biol Cybern* 87:440–5.

Sutton MA, Schuman EM (2006) Dendritic protein synthesis, synaptic plasticity, and memory. *Cell* 127:49–58.

Tanimoto H, Heisenberg M, Gerber B (2004) Experimental psychology: event timing turns punishment to reward. *Nature* 430:983.

Tzounopoulos T, Kim Y, Oertel D, Trussell LO (2004) Cell-specific, spike timing-dependent plasticities in the dorsal cochlear nucleus. *Nat Neurosci* 7:719–725.

Ultsch A, Schuster CM, Laube B, Betz H, Schmitt B (1993) Glutamate receptors of drosophila melanogaster. primary structure of a putative nmda receptor protein expressed in the head of the adult fly. *FEBS Lett* 324:171–7.

Urban NN, Sakmann B (2002) Reciprocal intraglomerular excitation and intra- and interglomerular lateral inhibition between mouse olfactory bulb mitral cells. *J Physiol* 542:355–367.

Vassar R, Chao S, Sitcheran R, Nunez J, Vosshall L, Axel R (1994) Topographic organization of sensory projections to the olfactory bulb. *Cell* 79:981–91.

Vassar R, Ngai J, Axel R (1993) Spatial segregation of odorant receptor expression in the mammalian olfactory epithelium. *Cell* 74:309–318.

Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287:1273–6.

Vogels TP, Rajan K, Abbott LF (2005) Neural network dynamics. *Annu Rev Neurosci* 28:357–76.

von der Malsburg C, Schneider W (1986) A neural cocktail-party processor. *Biol Cybern* 54:29–40.

Vosshall LB, Amrein H, Morozov PS, Rzhetsky A, Axel R (1999) A spatial map of olfactory receptor expression in the drosophila antenna. *Cell* 96:725–36.

Vosshall LB, Wong AM, Axel R (2000) An olfactory sensory map in the fly brain. *Cell* 102:147–59.

Wachowiak M, Cohen LB (1999) Presynaptic Inhibition of Primary Olfactory Afferents Mediated by Different Mechanisms in Lobster and Turtle. *J. Neurosci.* 19:8808–8817.

Wachowiak M, McGann JP, Heyward PM, Shao Z, Puche AC, Shipley MT (2005) Inhibition of Olfactory Receptor Neuron Input to Olfactory Bulb Glomeruli Mediated by Suppression of Presynaptic Calcium Influx. *J Neurophysiol* 94:2700–2712.

Wang JW, Wong AM, Flores J, Vosshall LB, Axel R (2003) Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. *Cell* 112:271–282.

Ward B, McGuinness L, Akerman CJ, Fine A, Bliss TVP, Emptage NJ (2006) State-dependent mechanisms of ltp expression revealed by optical quantal analysis. *Neuron* 52:649–661.

Wehr M, Laurent G (1996) Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384:162–6.

Wehr M, Laurent G (1999) Relationship between afferent and central temporal patterns in the locust olfactory system. *J Neurosci* 19:381–90.

Weiss M (1981) Structural patterns in the corpora pedunculata of orthoptera: a reduced silver analysis. *J. Comp. Neurol.* 203:515–553.

Wellis D, Scott J, Harrison T (1989) Discrimination among odorants by single neurons of the rat olfactory bulb. *J. Neurophysiolq* 61:1161–77.

Weth F, Nadler W, Korsching S (1996) Nested expression domains for odorant receptors in zebrafish olfactory epithelium. *Proc. Natl. Acad. Sci. USA* 93:13321–6.

Wicklein M, Strausfeld N (2000) Organization and significance of neurons that detect change of visual depth in the hawk moth manduca sexta. *J . Comp. Neurol.* 424:356–76.

Willmore B, Tolhurst D (2001) Characterizing the sparseness of neural codes. *Network: Comput. Neural Syst.* 12:255–70.

Wilson RI, Laurent G (2005) Role of gabaergic inhibition in shaping odor-evoked spatiotemporal patterns in the drosophila antennal lobe. *J Neurosci* 25:9069–79.

Wilson RI, Turner GC, Laurent G (2004) Transformation of olfactory representations in the drosophila antennal lobe. *Science* 303:366–70.

Wilson RI, Mainen ZF (2006) Early events in olfactory processing. *Annual Review of Neuroscience* 29:163–201.

Wolters A, Sandbrink F, Schlottmann A, Kunesch E, Stefan K, Cohen LG, Benecke R, Classen J (2003) A Temporally Asymmetric Hebbian Rule Governing Plasticity in the Human Motor Cortex. *J Neurophysiol* 89:2339–2345.

Wolters A, Schmidt A, Schramm A, Zeller D, Naumann M, Kunesch E, Benecke R, Reiners K, Classen J (2005) Timing-dependent plasticity in human primary somatosensory cortex. *J Physiol* 565:1039–1052.

Yao H, Dan Y (2001) Stimulus timing-dependent plasticity in cortical processing of orientation. *Neuron* 32:315–23.

Yokoi M, Mori K, Nakanishi S (1995) Refinement of odor molecule tuning by dendrodendritic synaptic inhibition in the olfactory bulb. *Proc Natl Acad Sci U S A* 92:3371–5.

Yu D, Keene AC, Srivatsan A, Waddell S, Davis RL (2005) Drosophila dpm neurons form a delayed and branch-specific memory trace after olfactory classical conditioning. *Cell* 123:945–57.

Zakharenko SS, Zablow L, Siegelbaum SA (2001) Visualization of changes in presynaptic function during long-term synaptic plasticity. *Nat Neurosci* 4:711–7.

Zars T, Fischer M, Schulz R, Heisenberg M (2000) Localization of a short-term memory in drosophila. *Science* 288:672–5.

Zhang LI, Tao HW, Holt CE, Harris WA, Poo Mm (1998) A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395:37–44.

Zhigulin VP, Rabinovich MI, Huerta R, Abarbanel HD (2003) Robustness and enhancement of neural synchronization by activity-dependent coupling. *Phys Rev E Stat Nonlin Soft Matter Phys* 67:021901.

Zou Q, Destexhe A (2007) Kinetic models of spike-timing dependent plasticity and their functional consequences in detecting correlations. *Biological cybernetics* 97:81–97.

Zou Z, Horowitz LF, Montmayeur JP, Snapper S, Buck LB (2001) Genetic tracing reveals a stereotyped sensory map in the olfactory cortex. *Nature* 414:173–9.